ISSUE 42, SPRING 2002
XCELL JOURNAL
XILINX, INC.

Issue 42
Spring 2002

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

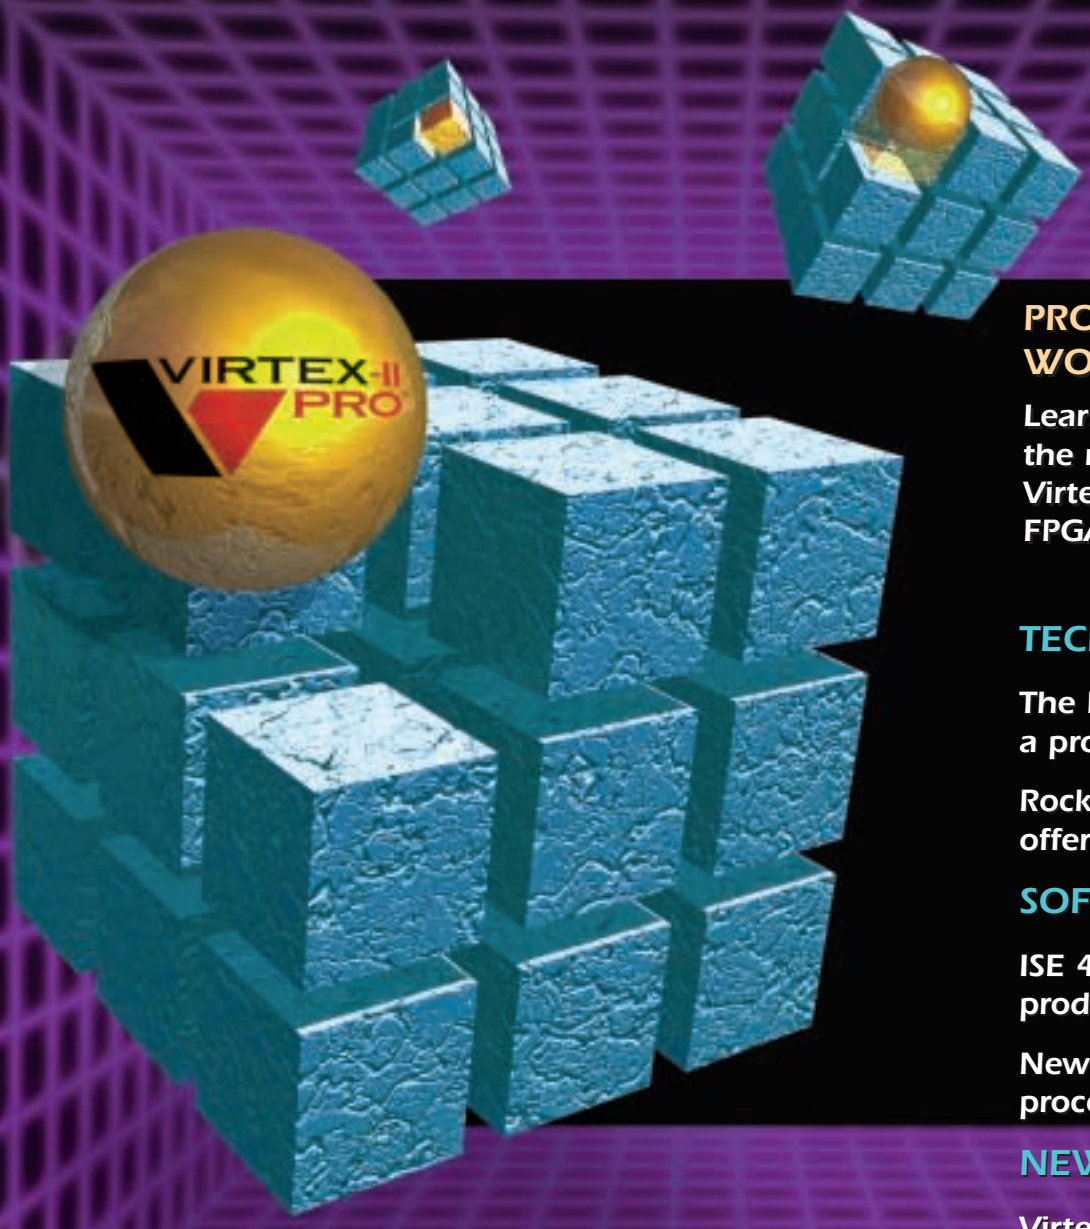

PROGRAMMABLE WORLD 2002

Learn all about the new Virtex-II Pro FPGAs

TECHNOLOGY

The PowerPC architecture: a programmer's view

Rocket I/O transceivers offer 3.125 Gbps capability

SOFTWARE

ISE 4.2i expands design productivity once again

New tools for embedded processor software design

NEWS

Virtex-II receives Product of the Year award

Cover Story

A revolutionary breakthrough in processing and system design, from Xilinx and IBM

XILINX®

Who Are You? What Did You Say?

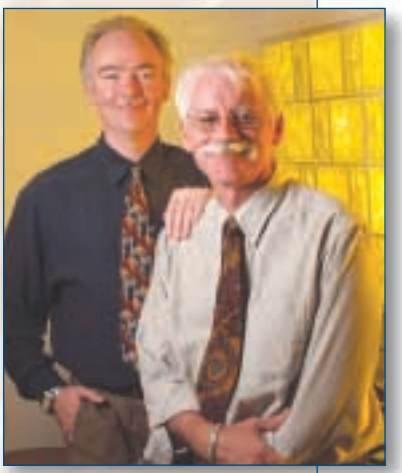

EDITOR IN CHIEF — Carlis Collins, editor@xilinx.com, 408-879-4519

MANAGING EDITOR — Tom Durkin, tom.durkin@xilinx.com, 530-271-0899

DESIGN & ILLUSTRATION — Scott Blair, Dan Teie

Xcell journal

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3450
Phone: 408-559-7778
FAX: 408-879-4780

Many of you have taken the time to give us your very valuable feedback about how we can continue to improve this *Xcell Journal*. After all, it is your journal, and its only purpose is to make your job easier and more productive, while also providing insights into the trends and technologies that are shaping the future of logic design. The overwhelming majority of responses indicated that Xcell is a huge success, often read cover to cover, and then saved for later reference. Thank you!

Here's some of what we learned from our reader survey:

- Most of you are design/development engineers (74%), doing digital logic design using FPGAs (88%) and CPLDs (76%), for industrial (38%), networking (35%), data processing (25%), and military (24%) applications, in companies of less than 500 employees (60%).
- Your three most popular categories are technical ("how to") articles, new product announcements, and the product reference guides. (You can bet these sections will grow fatter in future editions.) Your least popular category is Customer Success stories.
- Most of you read *Xcell Online* and would still like to receive the printed journal, monthly – almost everyone wanted more information, more often.

More Information, More Often

If you haven't noticed, *Xcell Online* is a fast and efficient way to stay informed about Xilinx and its partners. And, it keeps getting better. As we develop the website, you will see many more articles about the topics that are important to you. *And, you will find many valuable articles that we could not fit into the printed Xcell Journal* (we receive far more content than we can print). You can see *Xcell Online*, and other Xilinx literature, at: *www.xilinx.com/literature/*.

Virtex-II Pro and the New *Programmable World 2002*

This issue of the *Xcell Journal* focuses not only on a new product, but also a revolutionary new development paradigm. Our new Virtex-II Pro™ family of Platform FPGAs is truly unique and very powerful – it promises to change forever the way you approach system-level design. As this new technology quickly evolves, you will find the latest technical articles, online, as soon as we can publish them, on *Xcell Online*.

You will also have the opportunity to hear the industry leaders discussing this new design paradigm, and receive valuable training, at the Xilinx Programmable World 2002 event in April. Attend Programmable World 2002 either online or in person at various locations. See *www.xilinx.com/pw2002* for details.

Thanks again for your interest in Xilinx and Xcell. Our primary goal is to help you succeed, so please continue to send me your comments and suggestions.

Carlis Collins
editor@xilinx.com

View from the Top Page 4

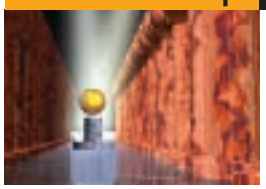

Interesting Times at Xilinx

Witness the worldwide, world-class debut of the Virtex-II Pro Platform FPGA solution for programmable systems.

Cover Story Page 6

The New Era of Programmable Systems

The next breakthrough in processing and system design methodology from Xilinx and IBM.

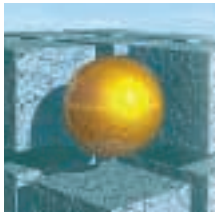

Technical Focus Page 10

Virtex-II Pro FPGAs: The Platform for Programmable Systems has Arrived

The Virtex-II Pro solution heralds a paradigm shift in system architecture by moving design based on zones of programmability to entire system-level programmability.

Technical Focus Page 16

The PowerPC Architecture: A Programmer's View

An introduction to the PowerPC programming model.

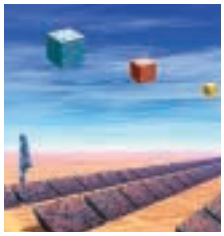

Xilinx News Page 44

SignOnce and Break the IP License Barrier

Xilinx sponsors the Common License Consortium to streamline the IP licensing process and improve your time to market.

Contents Spring 2002

Xcell journal

Interesting Times at Xilinx

Witness the worldwide, world-class debut of the Virtex-II Pro Platform FPGA solution for programmable systems.

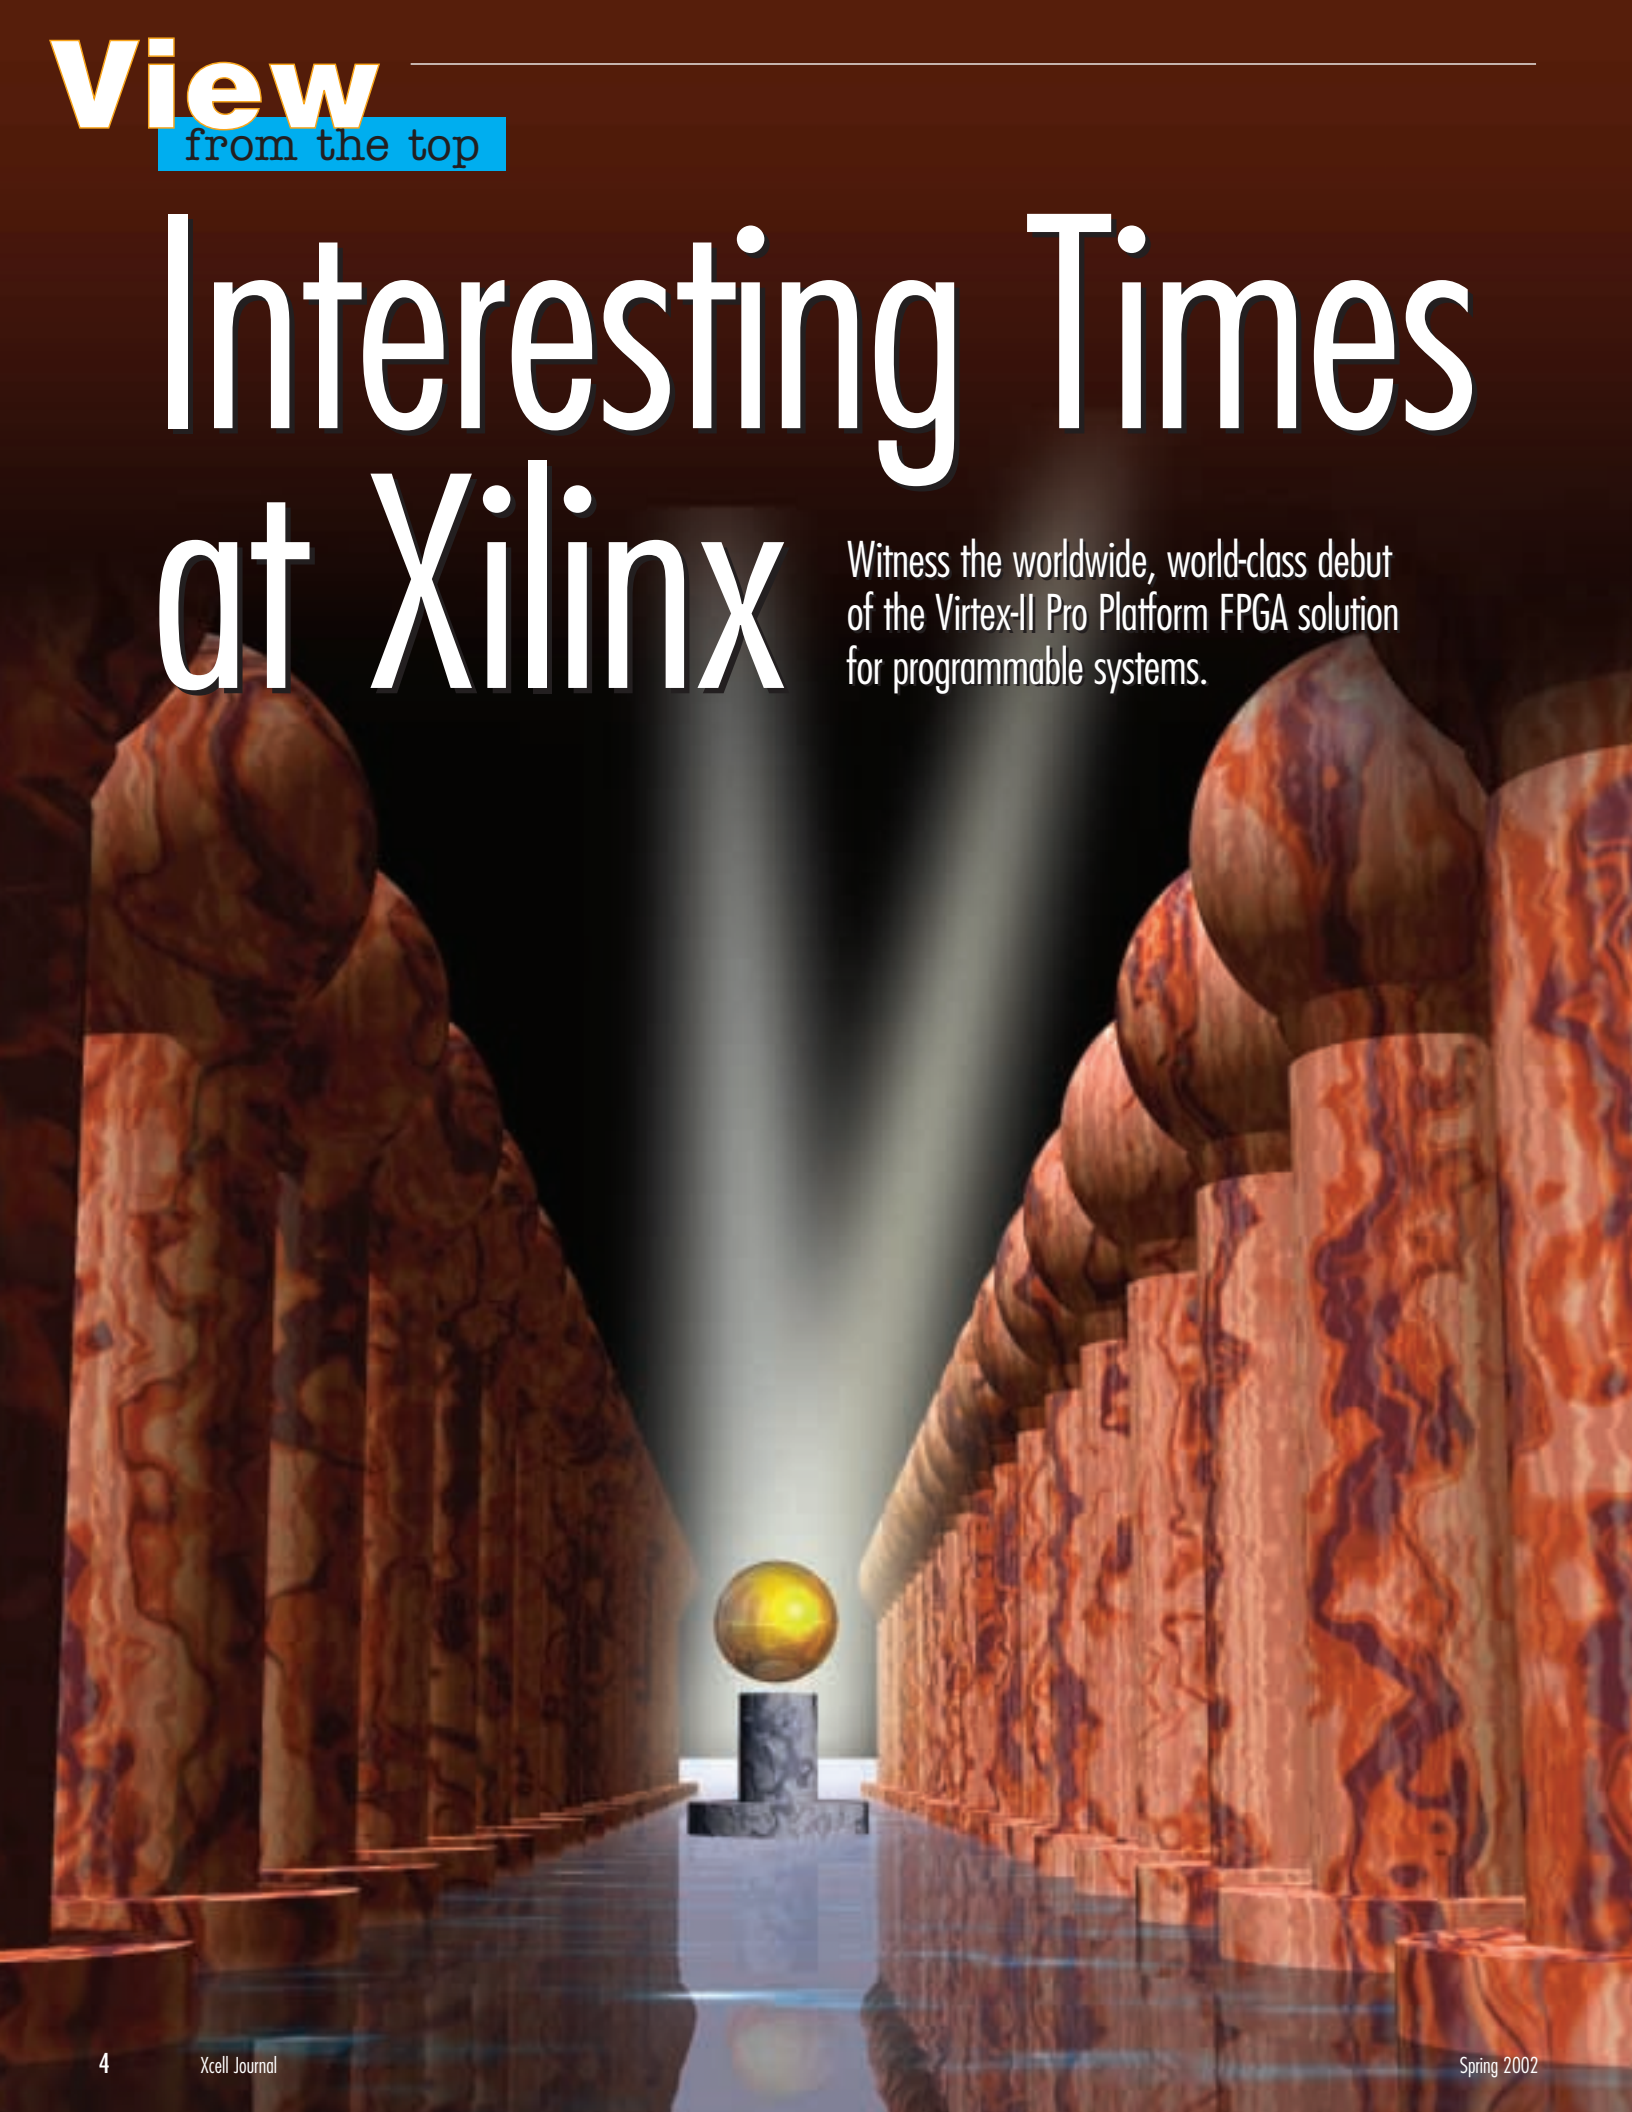

by Wim Roelandts
CEO, Xilinx

There is an ancient curse: May you live in interesting times. No doubt, we are living in interesting times. I choose, however, to view this as a blessing rather than a curse. True, competition is fierce and economic conditions are chaotic, but opportunities for growth and change in the programmable logic industry are limitless.

Big drops in revenues and even the sudden disappearance of large customers have forced us to revisit our strategies and redouble our efforts. Preserving capital, maintaining time to market, coping with lower budgets, reducing unnecessary risks, and staying the course in the face of uncertain market conditions has tested the mettle of all of us.

As we climb back up from the bottom of this recession, we are bringing with us a new paradigm for Xilinx®. With the introduction of the Virtex-II Pro™ Platform FPGA, we have changed from a programmable logic supplier to a purveyor of programmable systems.

What If?

Consider this:

- What if your corporation had access to an off-the-shelf, system-level product that allowed your design teams the maximum flexibility at system level without the traditional inventory risks?
- What if this off-the-shelf product had all the latest functionality they were looking for – and were way ahead of standard cell technologies?
- What if they could chose to never again deal with 0.13 micron silicon design issues or budget for huge NRE expenses?

The what-ifs for this dream can go on and on – but this is not a dream. This is what a Virtex-II Pro Platform FPGA solution can do for you right now.

Strategic Partnerships

Xilinx, IBM®, and Conexant™ Systems have been quietly working together to respond to the issues and challenges facing design teams and their corporations. Virtex-II Pro Platform FPGAs feature as many as four IBM PowerPC™ 405 processors immersed and embedded within the FPGA fabric. Moreover, the Virtex-II Pro devices connect to the outside

world via as many as 16 Rocket I/O™ 3.125 multi-gigabit serial transceivers capable of interfacing with multiple parallel and serial protocols and standards.

The Virtex-II Pro solution delivers both high-performance processing and high-bandwidth connectivity all in one device. And that's not all. Xilinx XCITE™ digitally controlled impedance technology removes hundreds of termination resistors from the printed circuit board. Xilinx IP Immersion and Active Interconnect technologies allow the PowerPC processors to bypass peripheral bus bottlenecks to connect directly with the FPGA logic and memory array.

Our partnerships on the software side with Wind River Systems, The MathWorks, Cadence Systems, Mentor Graphics, Synopsys, Synplicity, and more have paid off as well. The Virtex-II Pro solution comes with a complete set of Xilinx-specific embedded software tools for development, simulation, and debugging.

The close alliance with our partners and the tight integration of hardware and software in the Virtex-II Pro platform allows on-demand architectural synthesis with tremendous flexibility and scalability. You can efficiently divide complex functions between high-speed implementation in hardware and high-flexibility implementation in software.

See for Yourself

In the process of delivering the Virtex-II Pro solution, Xilinx has had to change its infrastructure to go beyond being a programmable logic supplier into becoming a programmable system provider. We knew that being a system-level provider didn't mean just innovation in silicon. It meant acquisitions and alliances in areas of I/O speed and connectivity, software development, design services – and preparing our entire workforce – from the experts in the field to the experts in customer support – to truly deliver a complete solution to our partners and customers.

Talk is cheap. So, let us show you what the family of Virtex-II Pro Platform FPGAs can do for you. I personally invite you and your engineering design teams to Programmable World 2002 to be held April 17 in San Jose, Boston, Paris – and more sites to be announced later. This worldwide exposition and conference will offer a general session and feature 16 presentations in four technical tracks.

Registration is mandatory, seating is limited, but participation is free. We, and our world-class partners, want to train you for the next generation of embedded, system-level programmable devices. For more information, read this issue of *Xcell Journal* (including the back cover) and register online at *www.xilinx.com/pw2002*.

The New Era of Programmable Systems

The next breakthrough in processing and system design methodology comes from the merger of the most advanced technologies from Xilinx and IBM.

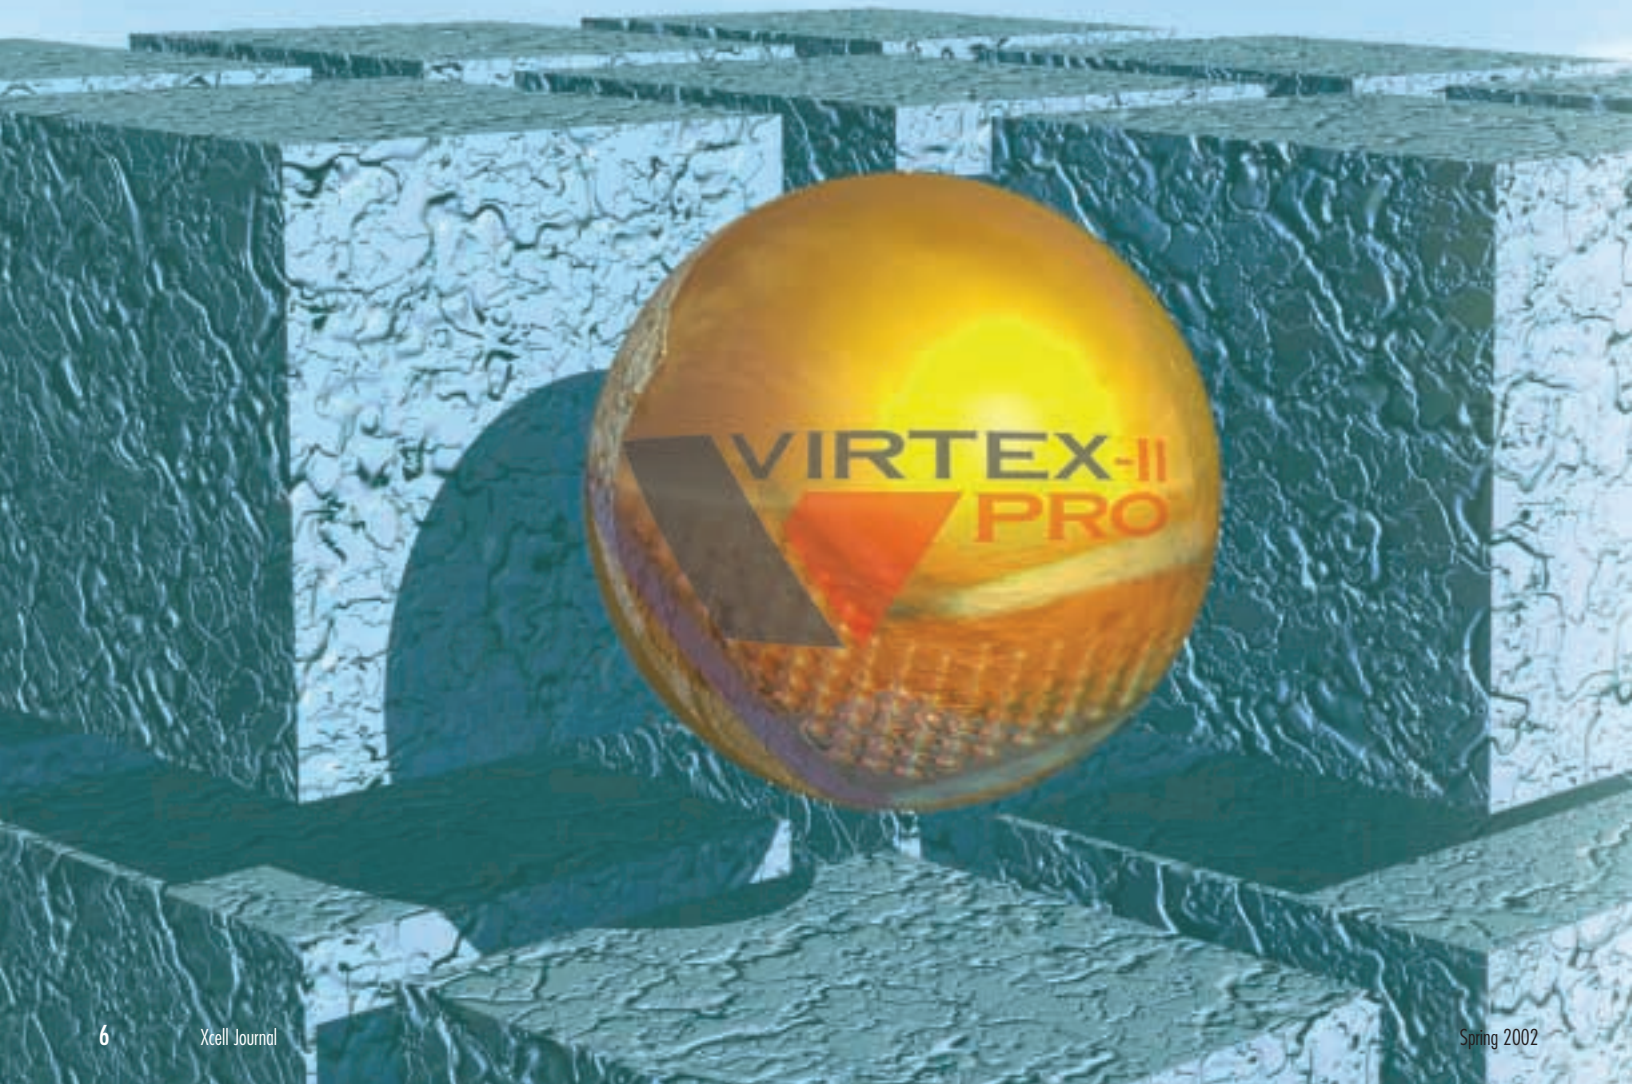

by Babak Hedayati
Sr. Director of Product Solutions and Partnerships
Xilinx, Inc.
babak.hedayati@xilinx.com

Over the course of the semiconductor revolution, with the help of Moore's Law, FPGAs have grown to densities of 10 million gates. They have consumed key system-level functions such as block memory, clock management, digitally controlled impedance matching, embedded multipliers, 844 Mbps LVDS I/Os, and many other functions. As the densities increased so did the insatiable hunger for soft IP cores for simple functions, complex DSP algorithms, networking protocols, interfaces, and so on.

The advantage gained by increased FPGA capabilities – sometimes unbelievable to some – has been adopted by thousands of design teams looking to improve their time to market by targeting segments of their system to Virtex™-II FPGAs whenever possible. Yet few could initially imagine the possibilities of a "programmable system" when Xilinx talked of immersing 300 MHz IBM® PowerPC™ 405 processors into the Virtex-II FPGA fabric and embedding high-speed multi-gigabit serial I/Os around it.

Many immediately saw the value of integration for cost reduction, increased performance, and reliability. Others saw the potential for its incredible flexibility and scalablity for implementing specialized and high-speed interfaces. The idea of extreme hardware parallel processing, and multiple processing on the same device, enticed many system designers to consider the great possibilities of such a solution.

Many engineers expected the next breakthrough in processing and system design methodology to come solely from the masters of the microprocessor world and leading ASIC vendors – but the breakthrough has come from the merger of the most advanced technologies from Xilinx and IBM.

The need for high-speed communication and increased bandwidth has driven the rapid evolution of technology throughout multiple industries. Design challenges associated with integration, high speed interfacing, higher performance processing, and new design methodologies must be solved, and the rapid rate of change in technology demands hardware programmability – this time at the system level.

Digital Convergence

The convergence of voice, video, data processing, and packet processing both on the infrastructure equipment and consumer products is putting immense pressure on corporations and their engineers. They must now incorporate computing, networking, wireless, and video imaging technologies that previously existed stand-alone in their respective markets.

This digital convergence has been an inevitable reality since the early days of the electronics industry, where specialized equipment and devices demonstrated their true potential as soon as they were connected with other devices. First came the telegraph, then the telephone, computers, video, the Internet, storage, wireless, and the infrastructure behind it all. Now the world is incredibly crowded with new consumer technologies that seem to pop up on a daily basis incorporating new features from the digital revolution. They target individual niche areas, compete, and often become extinct, just as suddenly as they were brought to market. This extinction is often caused by rapidly changing standards and requirements, or by competitive products putting tremendous pressure on corporations and their engineering teams.

Moving Toward Total Cost Management

Rapid product extinction makes executives question why their companies are spending so much capital on multi-million dollar ASIC NREs (Non-Recurring Engineering charges) and design automation contracts – because today's economic and technological conditions often require design changes midway in the development process. Success in today's marketplace is accomplished by getting to market first, not by designing for high volumes and the lowest unit cost – spending immense amounts in advance on creating custom ASICS, without a guaranteed future, is a recipe for failure.

Many companies are faced with huge inventories of ASICs and ASSPs that cannot be

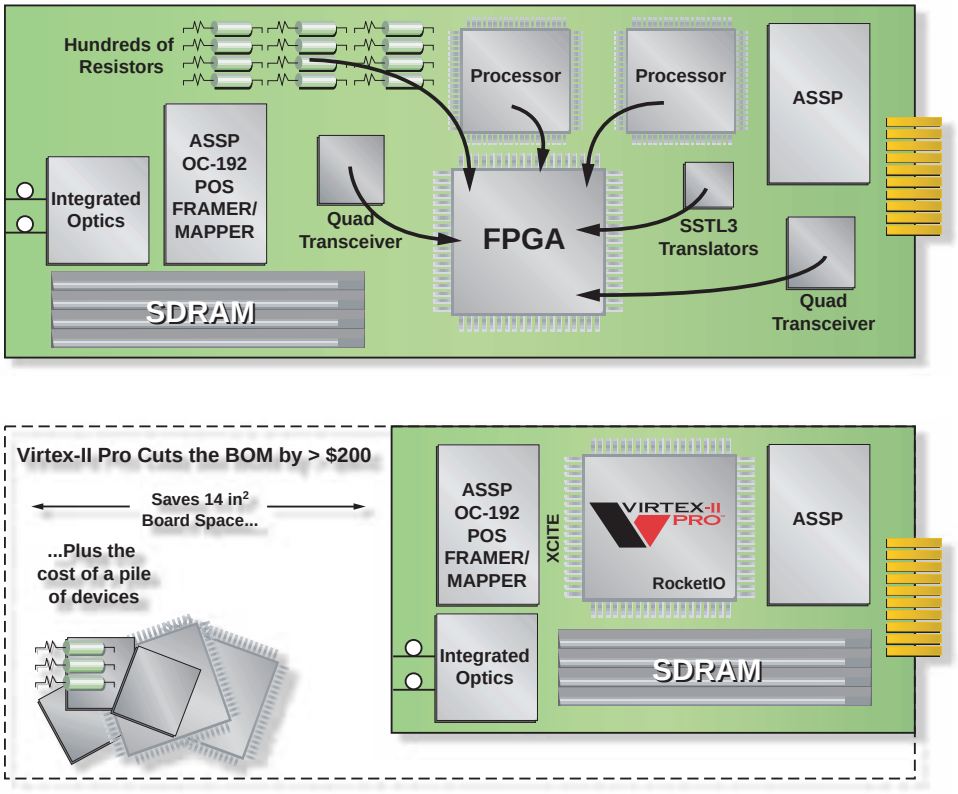

Figure 1 - Leading edge price/performance through integration and reduction of board size

re-targeted to multiple products to reduce the risk. Though ASICs are often less expensive for high volume designs, they incur much higher overall costs, higher risk, longer development cycles, and less time in market. It's time for corporations to evaluate the total cost of their development strategy as a whole and avoid falling into the pitfalls of separating capital investments, development costs, production costs, obsolescence costs, and inventory management.

The coming generation of computers and telecommunications equipment is very different from prior evolutions of information technology because it is dramatically reversing the age-old wisdom of creating specific devices for each application. Current systems are integrating computers, cell phones, game systems, cameras, appliances, automobiles, offices, and homes. Eventually, we will likely have only a few types of super systems remaining that synthesize and extend the capabilities of all current systems.

One of the key trends to reduce cost, increase performance, and increase the reliability of systems has been through integration. However when designers integrated their systems into custom ASICs, they increase inventory risks and require large initial up front investments. Hence, companies find it difficult to stop midway to change their designs. In addition, smaller companies or start ups find that they must commit the majority of their funding just to develop their platform – and sometimes they have a difficult time getting a large ASIC supplier to entertain their development.

A fully integrated system-level solution such as the Virtex-II Pro™ family solves all of these problems. Offering multiple gigabit serial I/Os, the fastest FPGA solution in the world, up to four Power PC 405s, XCITE™ controlled impedance technology, and other system-level features, you get a smaller board size, lower overall costs, and faster time to market, as illustrated in Figure 1.

The Ultimate Connectivity Solution

For a long time, the original PCI bus was the industry standard. To increase the bandwidth, many designers began to use bridges, continuing with the parallel, shared bus strategy. Then the standard moved to 64-bit 66 MHz versions, and later to PCI-X running at 133 MHz.

The problems with continuing this strategy are obvious. Wider busses require more pins and higher cost, and moving to higher frequencies causes signal integrity issues. Plus, the shared bus created more overhead and less bandwidth predictability. Although PCI will be used for years to come, today's performance-hungry applications have already moved toward packet switched LVDS-based parallel methodologies such as POS PHY Level 4, Flexbus 4, RapidIO, Hyper Transport, and others. This requires smarter protocols and point-to-point interfacing between devices and boards. The move was welcomed because it increased bandwidth. However, it often resulted in increased clock skew and signal integrity issues.

In the last two years companies like Xilinx, in partnership with Conexant (and the later acquisitions of companies such as RocketChips), have discovered how to implement mutli-gigabit serial I/Os in CMOS technology, making it a cost effective method of delivering point-to-point serial switched interconnections. This means much higher performance without any side effects. Other companies in the silicon industry, such as ASSP companies and standards committees, are now quickly adopting this strategy to reduce cost, increase reliability, and increase bandwidth, as shown in Figure 2 and Figure 3.

The Processing Revolution

The quest for higher performance processing is evident in many applications. Companies have traditionally turned to "farms" of expensive high-performance processors to achieve the performance they need. In doing so, they have usually faced prohibitive costs along with the massive efforts of managing and partitioning their tasks throughout the processor farm.

Hardware oriented companies such as networking, telecom, and wireless infrastructure developers have taken the lead by implementing parallel processing in hardware. For example, using Xilinx XtremeDSP™ solu-

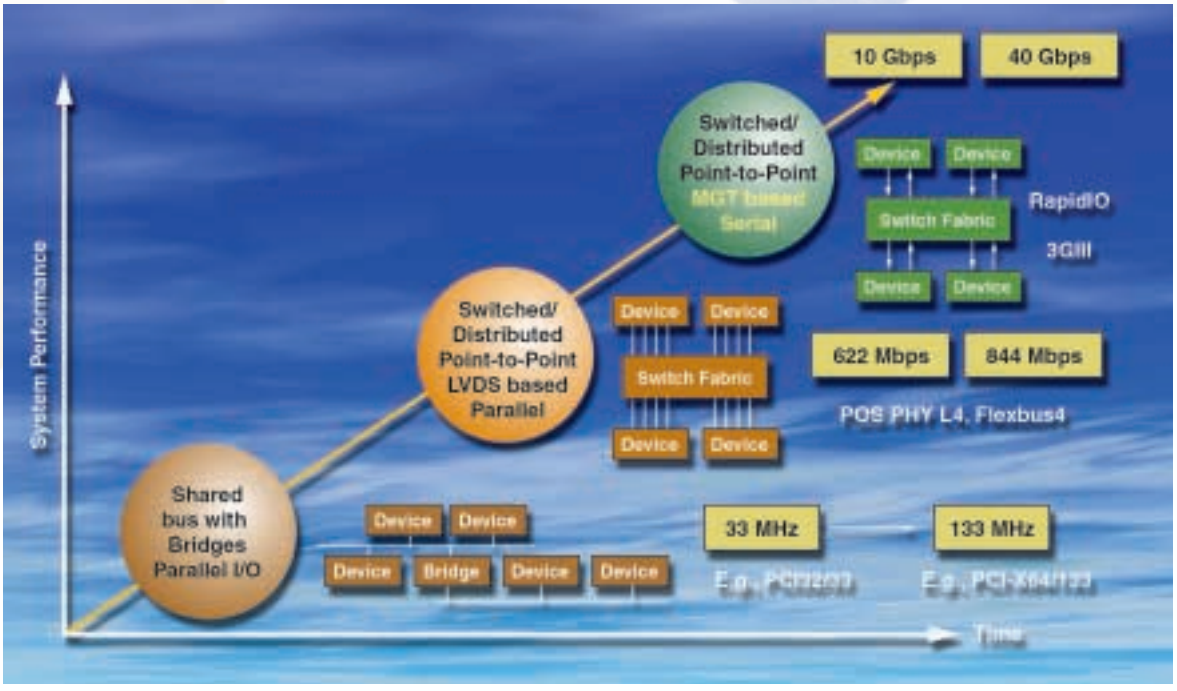

Figure 2 - The trend towards switched and distributed serial I/O

Figure 3 - Virtex-II Pro FPGAs support all connectivity standards

tions, with a single clock cycle they can process massive amounts of multiply accumulate functions (over 600 billion MACs per second). On the other hand, when companies have had to deal with multiple smaller tasks, they have traditionally turned to multiple processors, and optimized the code for processing each smaller task; a technique often used in network processors, as shown in Figure 4.

Now for the first time, through the Virtex-II Pro programmable system, designers get both high-performance processing and distributed processing through multiple PowerPC processors immersed in the FPGA fabric. With Virtex-II Pro FPGAs, the whole is much more than the sum of the parts.

Enabling a New System Development Paradigm

Design teams can now make system-level tradeoffs and optimization throughout the design cycle. Traditionally, architecture-level teams have had to make such tradeoffs early in system definition phase – deciding what goes into hardware what gets implemented as software code. This restriction has been the cause of many delayed products and products that have been unsuccessful, because of the inability to make adjustments for new features or performance optimizations during the design phase. Electronic design automation (EDA) companies have partially addressed the problem by developing system-level tools (such as behavioral partitioning tools and so on) to make the system-level tradeoffs easier.

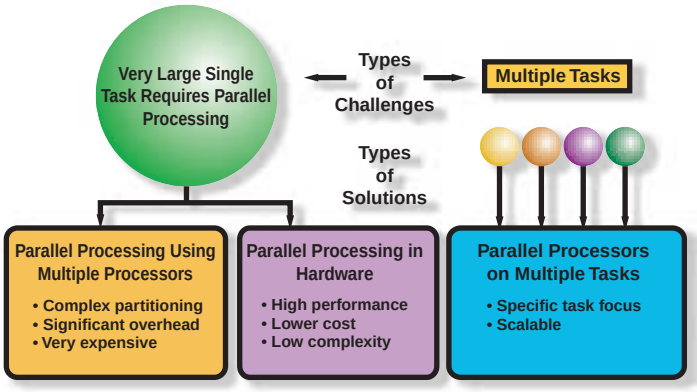

Figure 4 - Virtex-II Pro FPGAs provide higher performance processing

Now, with Xilinx technology, design teams can make tradeoffs and optimizations throughout the system design; creating a more integrated system with higher performance and faster time to market. They can even make changes to their hardware and software in the field, after the product is in the customers' hands, to fix bugs or implement new features. New business models can now be developed for programmable system design.

The Virtex-II Pro solution provides a standard programmable system platform, fully supported by embedded development tool vendors such as Wind River Systems; EDA companies such as Cadence, Mentor Graphics, and Synopsys; and system-level tools companies such as Celoxica and The Mathworks – all industry leaders and strategic partners to Xilinx. These partnerships enable a new development paradigm for programmable systems design. This overall solution of devices, software, cores, and partnerships means that companies like yours can now rest a little easier.

Welcome to the Programmable World

Xilinx and its partners are taking the next step in the evolution of programmable logic by creating a new event – Programmable World 2002. Here, you will hear industry leaders and visionaries discussing the latest Platform FPGA solutions, including detailed technical training for the PowerPC, Wind River tools, and others. From implementation techniques for multi-gigabit serial I/Os to digital signal processing, you will hear experts from more than 50 companies discussing this revolution in logic design.

Programmable World 2002 will be held simultaneously in multiple locations throughout North America and Europe. You won't need to travel, but if you do, it's all free with breakfast and lunch included.

April 17th you can see it all. To get the full details, or to register for the technical sessions, go to: *www.xilinx.com/pw2002*. Be sure to reserve your seat now; attendance is limited.

Virtex-II Pro FPGAs: The Platform for Programmable Systems Has Arrived

The Virtex-II Pro solution heralds a paradigm shift in system architecture by moving design based on zones of programmability to entire system-level programmability.

by Anil Telikepalli
Marketing Manager, Worldwide Marketing
Xilinx, Inc.
anil.telikepalli@xilinx.com

The curtains have been raised! The Virtex-II Pro™ Platform FPGA solution – the most sophisticated silicon and software product ever – is now available for programmable system design. Programmable systems represent flexible and scalable systems that are programmable at the architectural level. The goal in developing the Virtex-II Pro FPGA was to revolutionize system architecture by tightly integrating hardware and software functions on a single platform with unprecedented flexibility and scalability.

To achieve that objective, circuit engineers and system architects from IBM, Mindspeed, and Xilinx worked together to develop this advanced Platform FPGA. At the same time, engineering teams from top embedded systems companies, including Wind River Systems and Celoxica, worked alongside Xilinx software teams to develop the systems software and IP solutions that bring a new methodology to system design.

The result is the first Platform FPGA solution capable of implementing ultra-high bandwidth SOC (system-on-a-chip) designs that were previously the exclusive domain of custom ASICs. The Virtex-II Pro presents all the advantages of ASICs – and still retains all the flexibility and low development cost of programmable logic devices. The Virtex-II Pro solution enables high performance programmable systems specifically in the areas of wired and wireless networking, storage systems, professional broadcast, embedded systems, and digital signal processing systems. The new Virtex-II Pro FPGAs come in five densities, seven packages, and 15 combinations.

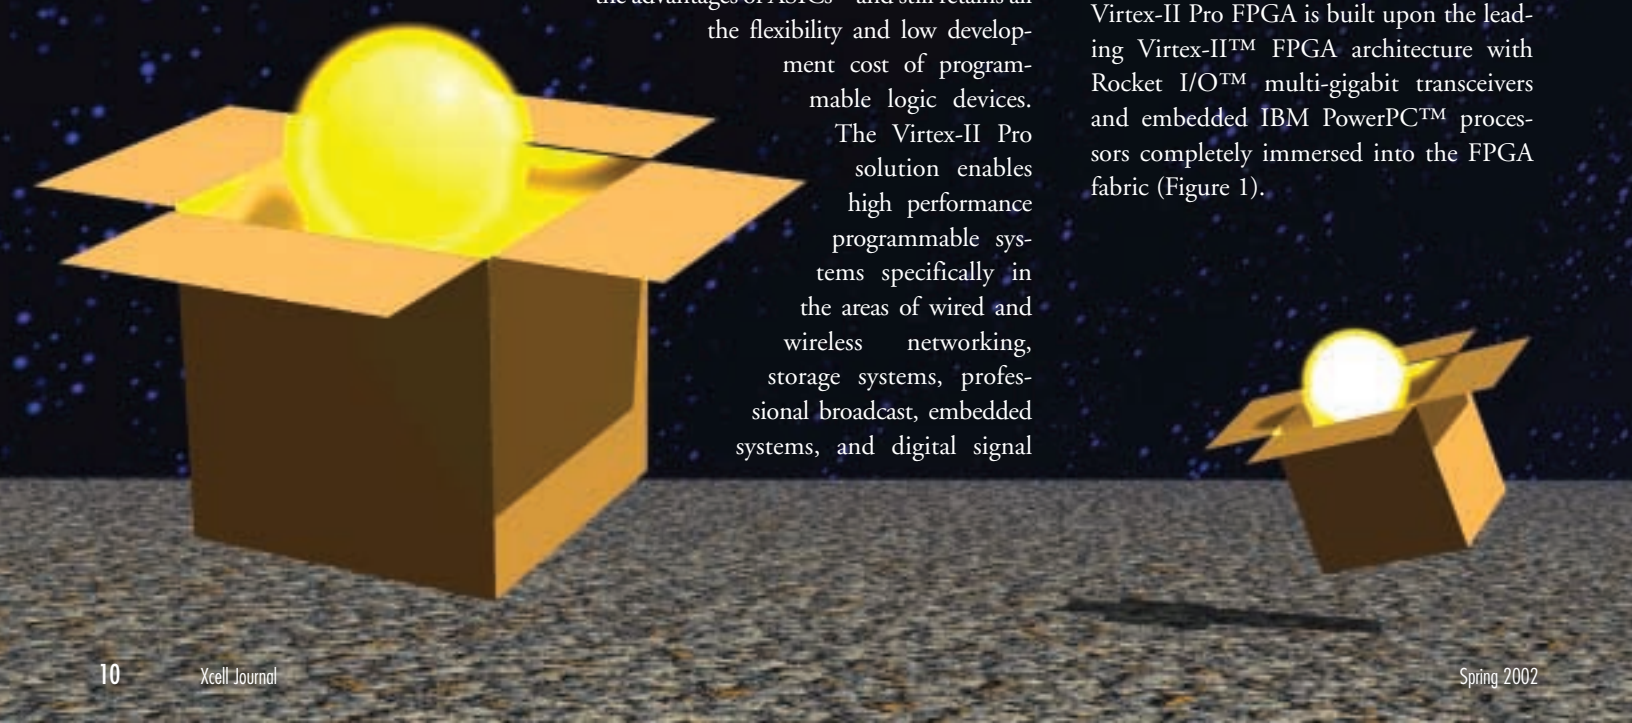

Virtex-II Pro FPGA Revealed

As a platform for programmable systems, the Virtex-II Pro FPGA is both flexible and scalable throughout all aspects of system architecture. By embedding processor cores within the FPGA fabric, the Virtex-II Pro architecture provides tight coupling between high-performance processors and the high-speed programmable logic. Together, the two components enable the most optimal yet flexible partitioning of hardware and software in a programmable system. The Virtex-II Pro FPGA is built upon the leading Virtex-II™ FPGA architecture with Rocket I/O™ multi-gigabit transceivers and embedded IBM PowerPC™ processors completely immersed into the FPGA fabric (Figure 1).

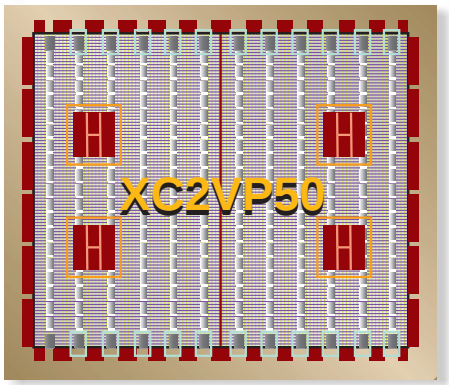

Figure 1 - The Virtex-II Pro XC2VP50 device features four IBM PowerPC 405 processors and 16 Rocket I/O multi-gigabit transceivers embedded in the FPGA fabric.

Additionally, Virtex-II Pro Platform FPGAs offer the following features:

- Five family members with 3,168 to 50,832 logic cells, and 216 Kb to 3,888 Kb of block RAM
- 0.13µ , 9-layer copper, low-k technology process
- 3.125 Gbps Rocket I/O multi-gigabit serial transceivers based on Mindspeed SkyRail™ technology, up to 16 per device
- 300+ MHz PowerPC embedded processor cores based on IBM's PowerPC 405 processor, up to four per device
- Virtex-II IP-Immersion technology powered by system-level features:
  - Flexible SelectI/O™-Ultra technology supporting 840 Mbps LVDS I/Os
  - Xilinx Controlled Impedance Technology (XCITE) capability, providing built-in digital impedance matching on all single-ended I/Os
  - Embedded 18 Kb dual-port block RAM resources
  - Embedded 18-bit x 18-bit multiplier blocks
  - DCM (digital clock manager) macros support de-skew and frequency/phase manipulation
  - Bitstream encryption (Triple-DES) for design protection.

Rocket I/O Transceivers

Rocket I/O multi-gigabit transceivers (MGTs) are based on Mindspeed SkyRail™ CMOS technology. Each full-duplex transceiver runs from 622 Mbps to 3.125 Gbps baud rate and includes the entire transceiver support circuitry (Figure 2). The Rocket I/O blocks are the first transceivers embedded in FPGAs to reach a baud rate of 3.125 Gbps. Up to 16 MGTs can be bonded together to provide an aggregate data rate of 40 Gbps for each of transmitter and receiver.

Well-designed serial transceivers have two fundamental requirements that distinguish them from others:

- Ability to operate at multi-gigabit rates to support emerging standards
- Ability to bundle multiple channels together for scalable data rate.

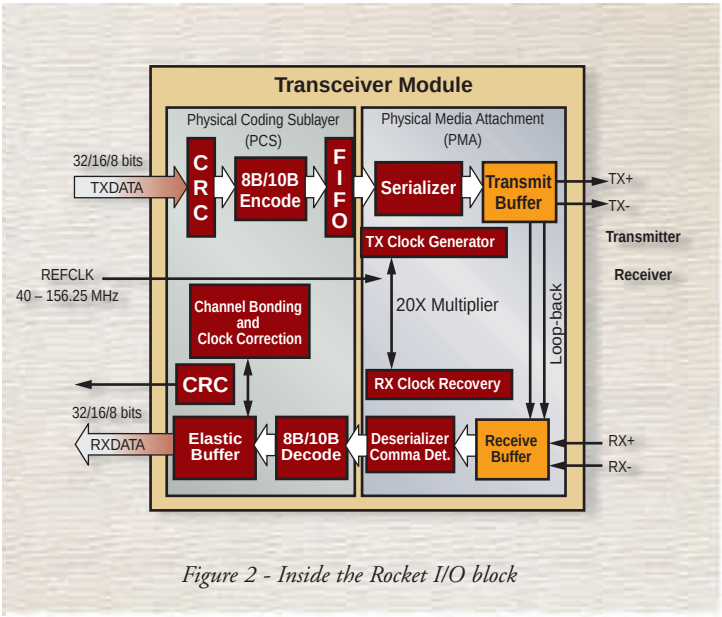

Figure 2 - Inside the Rocket I/O block

Each Rocket I/O transceiver consists of both a digital Physical Coding Sublayer as well as an analog Physical Media Attachment to provide a fully integrated serializer/deserializer function that enables the entire functionality and performance of emerging serial standards (Table 1).

Historically, serial transceivers have been analog components built using SiGe or GaAs processes. These transceivers generate enormous quantities of heat – and any integration of channels to increase the data rate was out of the question. However, Xilinx Rocket I/O transceivers not only provide multi-gigabit data rates, they can also be tied together to increase the aggregate bandwidth by using the built-in channel-bonding capability. This scalability is especially important as data rates increase and the industry moves toward designing compact optical networking equipment (impossible if you must put in several heat sinks).

For example, four Rocket I/O blocks allow 16 printed circuit board (PCB) traces to support full-duplex 10 Gbps data rates. This is equivalent to 256 traces of typical busses or 68 traces of a high-speed parallel bus. Thus, the four Rocket I/O blocks allow a 16X reduction of PCB traces over conventional parallel busses, resulting in significant reduction of PCB complexity and EMI system noise. In short, Rocket I/O technology allows higher bandwidth systems than currently possible, with cost savings from faster time-to-market, reduced power consump-

Technology	Line Speed
3GIO	2.5 Gbps
Serial ATA	1.5 Gbps
InfiniBand	2.5 Gbps
Gb Ethernet	1.25 Gbps
10 GE (XAUI)	3.125 Gbps
Serial RapidIO	1.25 Gbps
FibreChannel	1.06 Gbps

Table 1 - Virtex-II Pro Platform FPGAs support these protocols and baud rates.

tion, smaller PCB size, and lower component count (Figure 3).

Ultimate Connectivity

Virtex-II Pro Platform FPGAs, equipped with Rocket I/O MGTs, support emerging serial connectivity standards – and with Xilinx SelectI/O™-Ultra technology, these next-generation Platform FPGAs also support today's parallel connectivity standards (Table 2). Thus, the Virtex-II Pro FPGA serves as an ultimate connectivity platform to bridge across various interface standards in chip-to-chip, board-to-board, or even WAN/MAN/LAN networks.

In addition to these hardware physical interface capabilities, the Virtex-II Pro solution provides PowerPC processors and soft intellectual property (IP) cores to make designing with any protocol easy.

IBM PowerPC Processors

Each IBM PowerPC processor runs at 300+ MHz and 420 Dhrystone MIPS. Even though the PowerPC 405 core occupies a small portion of the die area, it provides tremendous system flexibility. Instead of attaching the PowerPC 405 processor next to the FPGA with a bus interface (as certain vendors have attempted), the Virtex-II Pro engineering team embedded the processor entirely within the FPGA fabric. Using Xilinx IP Immersion and Active Interconnect technologies, hundreds of processor nodes are directly connected to the FPGA logic and memory array.

Such total immersion gives you the utmost flexibility in hardware/software system architecture. You can efficiently divide complex functions between high-speed implementation in hardware and high-flexibility implementation in software. This direct-connect configuration bypasses the bottleneck of using a bus to interface between the FPGA and an attached/external processor.

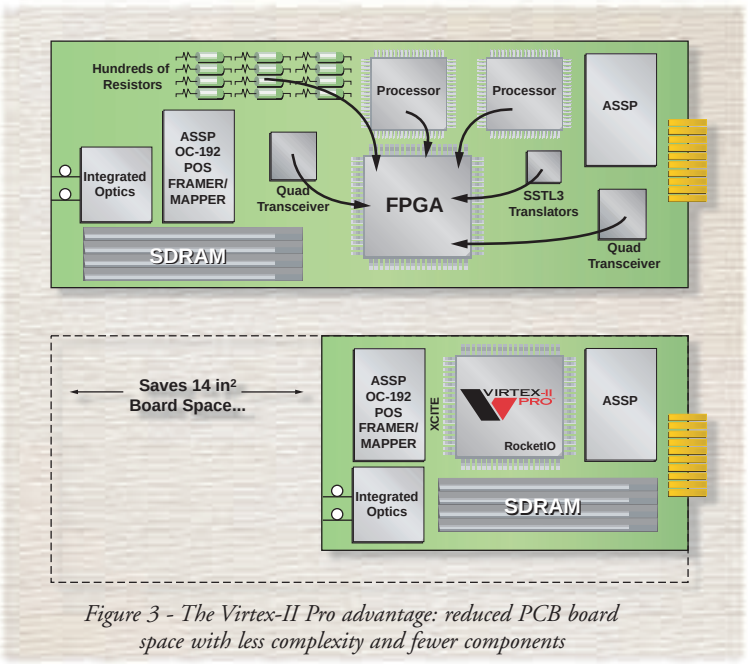

Figure 3 - The Virtex-II Pro advantage: reduced PCB board space with less complexity and fewer components

The PowerPC 405 core has unique on-chip memory (OCM) controllers that bypass the processor bus for fast, direct access to a fixed amount of instruction and data memory implemented in Xilinx SelectRAM™ modules (Figure 4). This is especially useful for data streaming applications.

The PowerPC processor is supported by IBM CoreConnect™ technology – a high-bandwidth 64-bit bus architecture that runs at 100 to 133 MHz. For maximum flexibility, the CoreConnect architecture is implemented as a soft IP within the Virtex-II Pro FPGA fabric (Figure 5). You can add CoreConnect peripherals from an extensive IP library from Xilinx and third-party partners or develop proprietary peripherals of your own.

The CoreConnect bus architecture has two main buses, called the Processor Local Bus (PLB) and the On-chip Peripheral Bus (OPB). These buses can be used for interfacing high-speed and low-speed peripherals with the PowerPC processor respectively. Additionally, a third Device Control Register (DCR) bus is used for transfers to and from general purpose peripheral device registers.

The Virtex-II Pro Platform FPGA is also supported by a complete set of embedded software tools for development and debug. Through an OEM agreement, Xilinx is able to provide software tools from Wind River Systems that are customized for Virtex-II Pro FPGAs. These include:

- Diab™ XE (Xilinx Edition) compiler
- SingleStep™ XE software debugger
- visionPROBE II XE JTAG run control hardware connection probe.

In addition, a suite of GNU (open-code Linux) tools is also available.

With the Virtex-II Pro solution, you can use the FPGA fabric for highly parallel processing and fixed algorithms. You can use the PowerPC processor for sequential com-

LAN/MAN/WAN	Board-to-Board		Chip-to-Chip	
● 10/100 Ethernet	● PCI 32/33	● RapidIO	● PCI 32/33	● POS-PHY L3/L4
● 1Gb Ethernet	● PCI 64/66	● CSIX	● PCI 64/66	● Flexbus 4
● 10Gb Ethernet	● PCI-X 100	● HyperTransport	● PCI-X66 & 100	● CSIX
● 1Gb Ethernet PHY	● Serial RapidIO	● InfiniBand	● RapidIO	● HyperTransport
● 10GE XAUI	● Fibre Channel	● 3GIO	● 10GE XAUI	● 3GIO
● SONET Standards*	● 10GE XAUI			

Table 2 - Virtex-II and Virtex-II Pro FPGAs offer multiprotocol connectivity.

● Serial standards enabled by Rocket I/O technology – Virtex-II Pro FPGAs
● Parallel standards enabled by SelectI/O-Ultra technology – Virtex-II & Virtex-II Pro FPGAs
* SONET compatible, supports data rate only

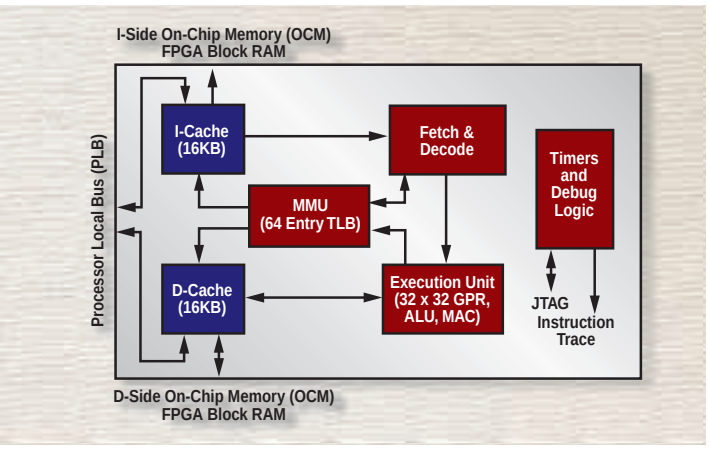

Figure 4 - Inside the PowerPC 405 processor

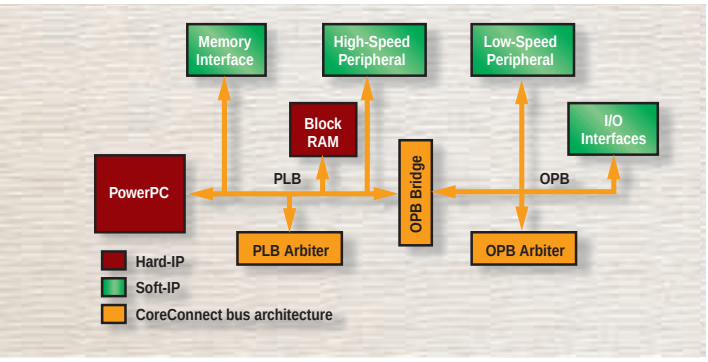

Figure 5 - Hard and soft IP cores

puting, exception handling, and control functions. The Rocket I/O serial transceivers and SelectI/O-Ultra parallel technologies enable optimal data access into and out of the Virtex-II Pro FPGA.

On-Demand Architectural Synthesis

Architectural synthesis is a combination of tools and technologies that allows designers to specify high-level requirements for designing their systems. In other words, architectural synthesis is a tool-based partitioning of hardware and software.

In order for architectural synthesis to work, both the hardware and software components of the system must be tightly integrated. Virtex-II Pro Platform FPGAs enable on-demand architectural synthesis with tremendously flexible, scalable, and high-bandwidth features. You can perform architectural synthesis anytime in the product cycle – during system design and debug phases, or even after the product has shipped. With abundant resources of hardware and software, Virtex-II Pro FPGAs give you the flexibility and scalability for fine-tuned system architecture, partitioned optimally between hardware and software.

Price & Performance Leader

Leading-edge systems need high bandwidth serial I/O, which has only been achievable by interfacing an FPGA to an external serial transceiver by means of hundreds of pins. Similarly, high-performance systems typically require one or more processors on the board, creating even more connectivity problems and PCB complexity.

By immersing multi-gigabit transceiver blocks and processor cores within the FPGA fabric (Figure 6), the Virtex-II Pro Platform FPGA delivers the best price and performance. Integrating processors and transceivers within the FPGA fabric lowers costs and raises performance by:

- Saving PCB space
- Simplifying PCB complexity
- Requiring fewer components
- Eliminating complex device interconnectivity issues
- Reducing overall system power consumption
- Using XCITE digitally controlled impedance technology to do away with external termination resistors
- Enabling optimal system partitioning between hardware and software.

A New Development Paradigm

By tightly integrating flexible and scalable high-performance programmable logic with the PowerPC processor, the Virtex-II Pro Platform FPGA fundamentally changes the way systems are designed. The Virtex-II Pro solution facilitates a new paradigm in system development with significant benefits, specifically in software engineering productivity:

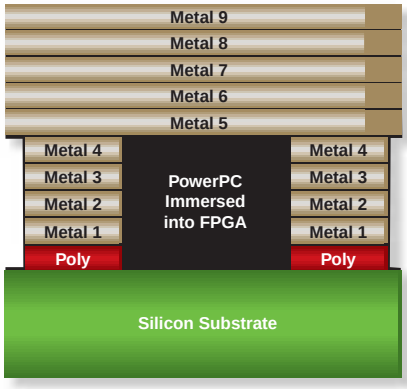

Figure 6 - Hard IP cores are deeply embedded and actively connected within the FPGA fabric

- Embedded processors can be used for rapid system pre-production to facilitate accelerated software development. A preliminary hardware platform can be built quickly by emulating a C-based algorithm using the embedded processors. This creation of a preliminary hardware platform allows software development to start much earlier in the design process, compared to current practice.
- Software debugging can be performed at hardware speeds while the hardware implementation continues to be speed-optimized. Xilinx ChipScope™ Pro on-chip verification tools provide in-system observability into both the FPGA hardware and the processor bus transactions.

Conclusion

The Virtex-II Pro Platform FPGA solution encompasses the following:

- Rocket I/O transceivers and IBM PowerPC processors immersed and embedded into the high performance Virtex-II Pro FPGA fabric
- Intellectual property solutions, including soft peripherals and connectivity cores
- Complete design resources, including development tools and kits

To find out more about this revolutionary next-generation Platform FPGA for programmable systems, go to *www.xilinx.com/virtex2pro.*

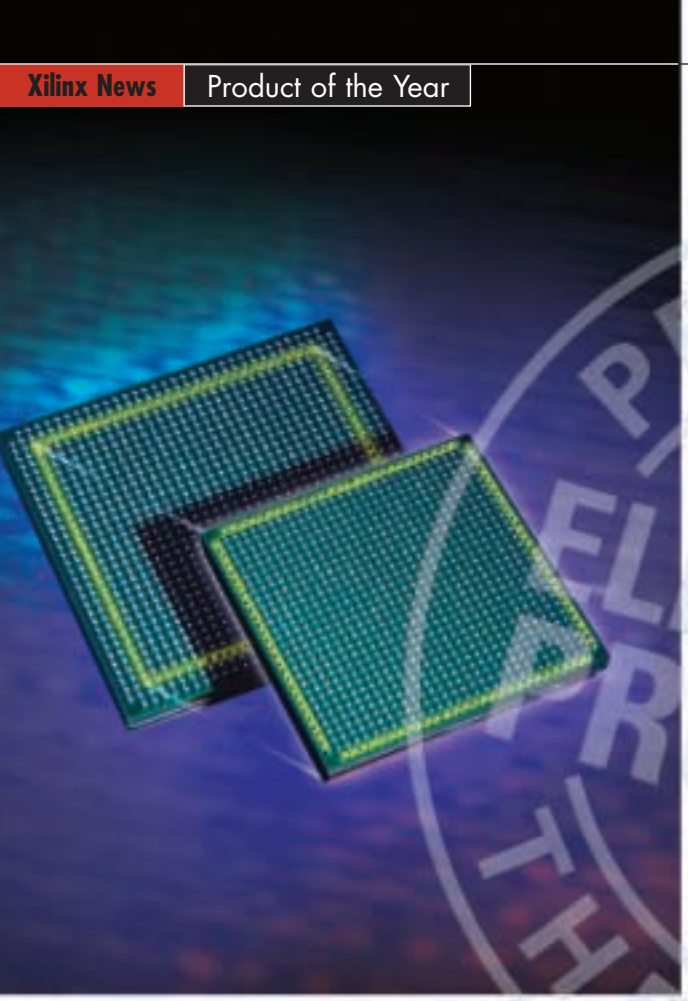

Virtex-II FPGAs — Product Of The Year

Xilinx recently received the Product of the Year Award from Electronic Products magazine — and we were the only programmable logic supplier to receive this award.

by Xilinx Staff

For 26 years, Electronic Products has held an annual contest to choose the most outstanding products introduced each year. The editorial board at the magazine considers thousands of product introductions based on significant advances in technology or its application, a decided innovation in design, or a substantial gain in price-performance benefits. As usual, picking winners was made difficult by the many impressive products announced during the year.

The Xilinx Virtex™-II FPGA family was recognized as the FPGA platform to address next-generation designs. "Programmable logic has heretofore been limited to relatively simple computational tasks and glue logic functions," said David Suchman, Digital IC editor at Electronic Products. "Now, for the first time, a programmable platform is available from Xilinx to enable rapid development of today's technically challenging applications."

The ever-increasing requirements for higher performance and system-level features are bringing new challenges as designers develop the next-generation of complex high-performance digital applications such as data communications and DSP systems. Characterized by high logic integration, fast and complex routing of wide buses, and extensive requirements for pipeline and FIFO memory, these new systems exceed the capabilities of current programmable logic devices, which lack the gate capacity, memory, routing resources, performance, and architecture flexibility that is required to fully support these designs. The Electronic Products' editors judged that the Xilinx Virtex-II series (see Electronic Products, May 2001) also solves the problems resulting from signal integrity, system timing, EMI, and security issues in these complex systems.

The Virtex-II family allows unlimited design changes throughout the development and production phases for optical networks, gigabit routers, wireless cellular base stations, modem arrays, and video broadcast systems. Capable of handling designs from 40,000 to 10 million system gates, Virtex Series FPGAs feature an interconnect architecture for optimizing routing, and an advanced memory array with up to 4.5 Mbits of on-chip memory. An additional feature is the industry's first digitally controlled impedance technology, which maintains constant impedance even with temperature and voltage fluctuations – eliminating hundreds of termination resistors, saving board space, increasing reliability, and lowering costs.

For more information on the winners, go to *www.electronicproducts.com*.

"Programmable logic has heretofore been limited to relatively simple computational tasks and glue logic functions. Now, for the first time, a programmable platform is available from Xilinx to enable rapid development of today's technically challenging applications." — David Suchman, Digital IC editor at Electronic Products.

DESIGN. TOGETHER.

At Cadence we offer a different perspective on the complex world of electronic design. It's an approach that focuses all our resources—from industry-leading technologies and services to strong partnerships and open collaboration—toward complete customer success.

Cadence Design Systems, Inc. • 2655 Seely Avenue • San Jose, CA 95134 • www.cadence.com

Visit us at the Embedded Systems Conference and Xilinx Programmable World 2002

The PowerPC Architecture:

A Programmer's View

An introduction to the PowerPC programming model.

by Anthony Marsala
IBM

The PowerPC™ Architecture is a Reduced Instruction Set Computer (RISC) architecture, with over two hundred defined instructions. PowerPC is RISC in that most instructions execute in a single-cycle and typically perform a single operation (such as loading storage to a register, or storing a register to memory). This article will focus solely on 32-bit implementations, which are the most widely available today.

The PowerPC architecture employs a layered approach, in that it is broken up into three levels or "books". By segmenting the architecture in this way, code compatibility can be maintained across implementations while leaving room for implementations to choose levels of complexity for price/performances tradeoffs. The three levels are broken up from the most general and common across implementations to the most operating system specific. The levels are:

- **Book 1.** ***User Instruction Set Architecture*** – This level defines the base set of instructions and registers that should be common to all PowerPC implementations.
- **Book 2.** ***Virtual Environment Architecture*** – This level defines additional user-level functionality that is outside the normal application software requirements. Areas include cache management, atomic operations, and user-level timer support.
- **Book 3.** ***Operating Environment Architecture*** – This level defines privileged operations typically required by an operating system. Areas include memory management, exception vector processing, privileged register access, and privileged timer access.

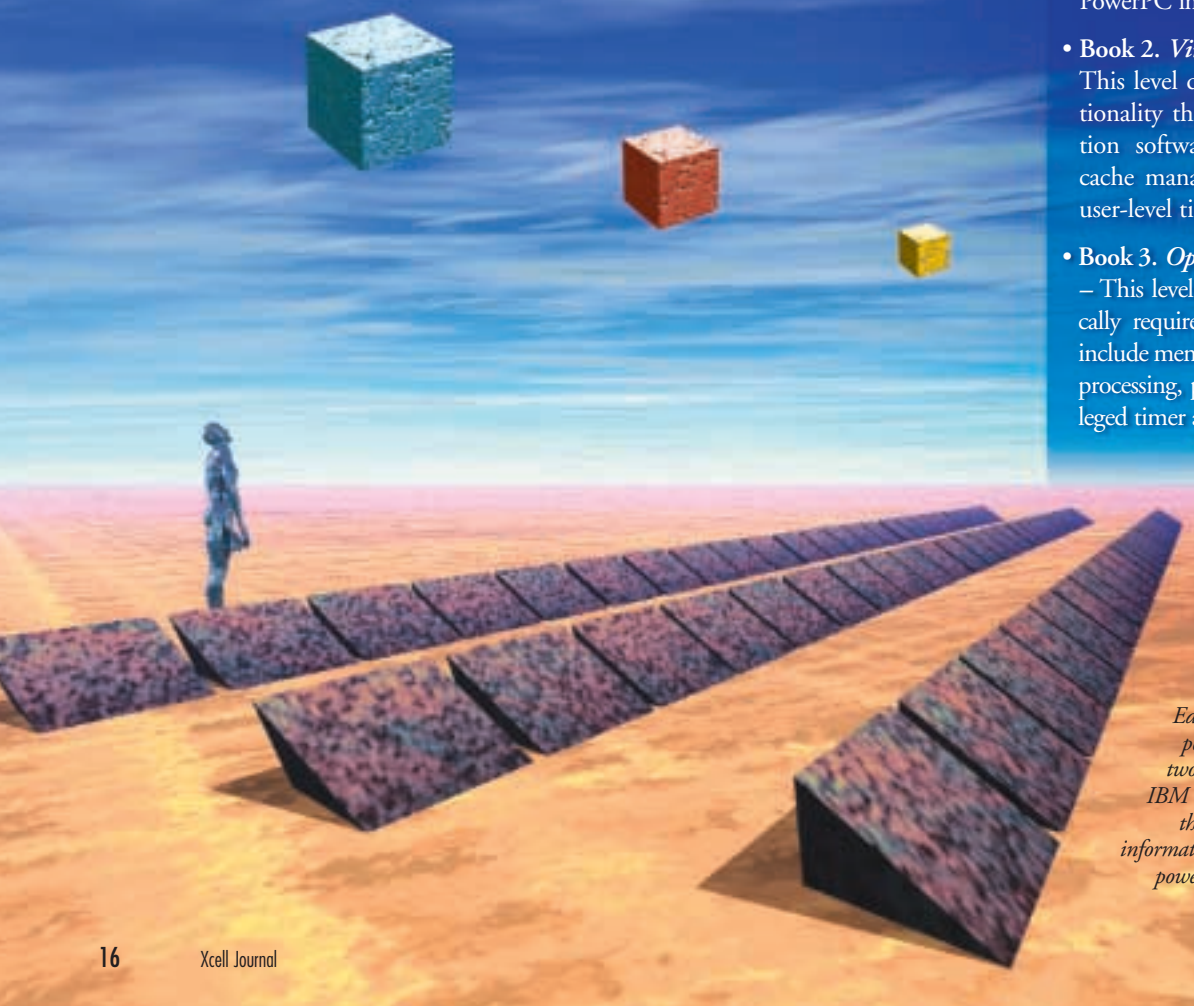

Editor's note: This article is reprinted with permission from IBM. It was originally a two-part series that ran in the April, 2001, IBM PowerPC Processor News. You can view the original articles, and find other useful information at: www-3.ibm.com/chips/products/powerpc/newsletter/apr2001/design-h-t.html.

Deviations from the original PowerPC Architecture offer flexibility to allow for enhancements that may come over time. In addition, IBM has defined its own Virtual Environment and Operating Environment levels for its PowerPC 400 family of embedded controllers.

Book E – A New Definition

A new PowerPC architecture update has been developed. Called "Book E", it combines the original three architecture levels into one new specification. This new specification also streamlines the definition of 64-bit implementations and eliminates non-substantive differences between IBM and Motorola implementations. The new standard maintains 100% code compatibility with Book 1 instructions and registers, while formally defining software-based memory management, a two-level interrupt hierarchy, and user-extendible instruction space for auxiliary processors. All of these enhancements address the needs of embedded systems.

To distinguish between the original architecture, the IBM embedded definitions, and Book E, the original architecture will be referred to as the "classic" architecture for the remainder of this article.

Storage Model

The 32-bit PowerPC architecture has native support for byte, halfword (16-bits), and word (32-bit) data types. Also, PowerPC implementations can handle string operations for multi-byte strings up to 128 bytes in length. The 32-bit PowerPC implementations support a 4 GB address space (2^{32}). All storage is byte addressable. For misaligned data accesses, alignment support varies by product family, with some taking exceptions and others handling the access through multiple operations in hardware.

Endianness

Classic PowerPC and the IBM PowerPC 400 family are primarily big-endian machines, meaning that for halfword and word accesses, the most-significant byte (MSB) is at the lowest address. Support for little endian varies by implementation. Classic PowerPC had minimal support, while the 400 family provides more robust support for little endian storage.

Book E is endian-neutral, as the Book E architecture fully supports both accessing method.

Registers

Classic PowerPC registers are broken into two classes: special-purpose registers (SPRs) and general-purpose registers (GPRs). IBM's PowerPC 400 family and Book E also define a third class of registers, called device control registers (DCRs), to address peripheral registers outside of the processor core in an embedded controller implementation. The three classes are explained below.

SPRs

SPRs give status and control of resources within the processor core. Table 1 shows different types of SPRs and their purpose. Where a single register exists, the SPR name is listed in parenthesis.

Supervisor vs User-Mode SPRs

When the processor is first initialized, it is in supervisor (also called privileged) mode. In this mode, all processor resources, including registers and instructions, are accessible. The processor can limit access to certain privileged registers and instructions by placing itself in user (also called problem-state) mode. This protection limits application code from being able to modify global and sensitive resources, such as the caches, memory management system, and timers. Mode switching is controlled via the Machine State Register.

- The Instruction Address Register (IAR) is known to programmers as the program counter or instruction pointer. It is the address of the current instruction. This is really a pseudo-register, as it is not directly available to the user. The IAR is primarily used by debuggers to show the next instruction to be executed.
- The Processor Version Register (PVR) is useful for code common across multiple processors that must make decisions based on a specific processor.

User-Mode SPRs

There are four SPRs available in user-mode that are important to understand:

- The Link Register (LR) is a 32-bit register that contains the address to return to at the end of a function call. Certain branch instructions can automatically load the LR to the instruction following the branch.

SPR Register Type	Access Mode	Purpose
Count (CTR)	User	Branching and Loop Control
Link (LR)	User	Subroutine Branching
Save/Restore	Supervisor	Interrupt Context Save
Debug	Supervisor	On-chip Debug Capabilities
Timers	User (read) Supervisor (write)	Timing Facilities
Interrupt Vector Prefix	Supervisor	Locates Interrupt Addresses
Exception	Supervisor	State information where exceptions occur
Storage Attribute Control	Supervisor	Controls Storage Attributes (W,I,G,LE)
Processor Version (PVR)	Supervisor	Identifies PowerPC Implementation
General Purpose (SPRGn)	Supervisor	Used by Operating Systems
Integer Exception (XER)	User	Carry Bit, Overflow, String Lengths
MMU	Supervisor	Instruction/Data Translation Control

Table 1: SPR registers

The **blr** instruction moves the program counter to the address in the LR.

- The Fixed Point Exception Register (XER) contains overflow information from fixed point arithmetic operations. It also contains carry input to arithmetic operations and the number of bytes to transfer during load and store string instructions **lswx** and **stswx**.
- The Count Register (CTR) contains a loop counter that is decremented on certain branch operations. Also, the conditional branch instruction **bcctrx** branches to the value in the CTR.
- The Condition Register (CR) is grouped into eight fields, where each field is 4 bits that signify the result of an instruction's operation: Equal (EQ), Greater Than (GT), Less Than (LT), and Summary Overflow (SO).

Machine State Register (MSR)

MSRs represent the state of the machine. It is accessed only in supervisor mode, and contains the settings for things such as memory translation, cache settings, interrupt enables, user/privileged state, and floating point availability. Exact control bits vary by implementation.

The MSR does not readily fit into the SPR/DCR/GPR classification, as it contains its own pair of instructions (**mfmsr / mtmsr**) to read and write the contents of the MSR into a GPR.

DCRs

DCRs are similar to SPRs in that they give status and control information, but DCRs are for resources outside the processor core. DCRs allow for memory-mapped I/O control without using up portions of the 32-bit memory address space.

GPRs

The User Instruction Set Architecture (Level 1) specifies that all implementations have 32 GPRs (GPR0 - GPR31). GPRs are the source and destination of all fixed-point operations and load/store operations. They also provide access to SPRs and DCRs. They are all available for use in every instruction with one exception: In certain instructions, GPR0 simply means "0" and no lookup is done for GPR0's contents.

Instructions

Table 2 lists different instruction categories, and the types of instructions that exist in that category.

Instruction Category	Base Instructions
Data Movement	load, store
Arithmetic	add, subtract, negate, multiply, divide
Logical	and, or, xor, nand, nor, xnor, sign extension, count leading zeros, andc, orc
Comparison	compare algebraic, compare logical, compare immediate
Branch	branch, branch conditional, branch to LR, branch to CTR
Condition	rand, crnor, crxnor, crxor, crandc, crorc, crnand, cror, cr move
Rotate/Shift	rotate, rotate and mask, shift left, shift right
Cache Control	invalidate, touch, zero, flush, store, dcread, icread
Interrupt Control	write to external interrupt enable bit, move to/from machine state register, return from interrupt, return from critical interrupt
Processor Management	system call, synchronize, eieio, move to/from device control registers, move to/from special purpose registers, mtcrf, mfcr, mtmsr, mfmsr
MMU Control	TLB search, TLB read, TLB write, TLB invalidate all, TLB synchronize
MAC Unit	multiply low/high halfword and accumulate/subtract

Table 2 - Instruction categories

AND	OR	Exclusive OR	Rotate and Mask	Shift	Misc.
and	or	xor	rlwimi	slw	cntlzw
and.	or.	xor.	rlwimi.	slw.	cntlzw.
andi.	ori	xori	rlwinm	sraw	
andis.	oris	xoris	rlwinm.	sraw.	extsb
			rlwnm	srawi	extsb.
nand	nor	egv	rlwnm	srawi.	
nand.	nor.	egv.		srw	extsh
				srw.	extsh.
andc	orc				
andc.	orc.				

Table 3 - Power PC logical instructions

Deciphering an Instruction

For 32-bit implementations, all instructions are 32 bits (4 bytes) in length. Bit numberings for PowerPC are opposite of most other definitions; bit 0 is the most significant bit, and bit 31 is the least significant bit. Instructions are first decoded by the upper 6 bits, in a field called the primary opcode. The remaining 26 bits contain operands and/or reserved fields. Operands can be registers or immediate values.

Arithmetic Instructions

Many instructions exist for performing arithmetic operations, including add, subtract, negation, compare, multiply and divide. Many forms exist for immediate values, overflow detection, and carry in and out. Multiply and divide

instruction performance varies among implementations, as these are typically multi-cycle instructions.

Logical Instructions

Table 3 lists PowerPC logical instructions. Looking at the AND instruction, The "i" form means that a 16-bit immediate is used for the AND, the "is" form means that a 16-bit immediate is used in the upper 16-bits of the AND. For all "." forms, the CR[CR0] is updated as previously described. PowerPC has the ability to perform a 32-bit rotate-and-combine with a mask in a single cycle. In the miscellaneous column are instructions to count the leading zeros in a register, and sign extension instructions.

Load/Store Instructions

All loads/stores are performed using the GPRs. Instructions exist for byte, halfword, and word sizes. Special instructions include:

- Multiple-word load/stores (**lmw / stmw**), which can operate on up to 31, 32-bit words
- String instructions, which can operate on up to 128-byte strings
- Memory Synchronization instructions **lwarx** (Load Word and Reserve Indexed) and **stwcx.** (Store Word Conditional Index) are used to implement memory synchronization. lwarx performs a load and sets a reservation bit internal to the processor and hidden from the programming model. The associated store instruction **stwcx.** performs a conditional store only if the reservation bit is set and thereafter clears the reservation bit. CR[CR0]EQ is set to the state of the reservation bit at the start of the instruction so that software can determine if the write was successful.

Synchronization Instructions

Commonly misunderstood PowerPC instructions are those that perform synchronization. These instructions include:

- Enforce In/Order Execution of I/O (eieio) – This instruction is for data accesses to guarantee that loads and stores complete with respect to one another. Since PowerPC defines a weakly ordered storage model in which loads and stores can complete out of order, this instruction exists to guarantee ordering where necessary.

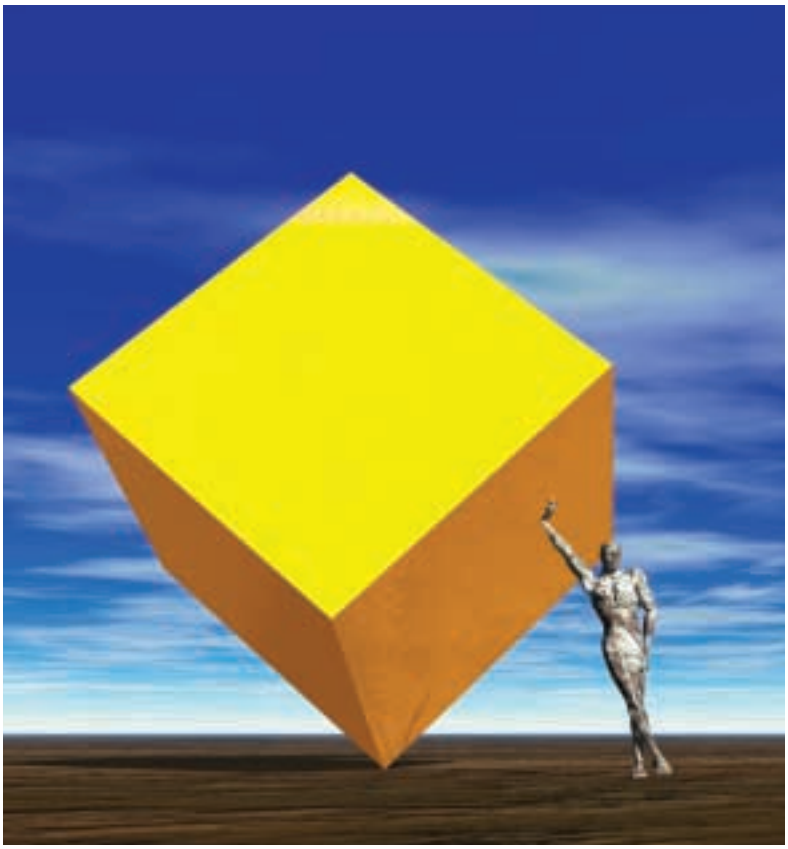

- Synchronize (sync) – This instruction guarantees that the preceding instructions complete before the sync completes. This instruction is useful for guaranteeing load/store access completion. For example, a sync may be used when writing memory mapped I/O registers to a slow device before making further access to the device.
- Instruction Synchronize (isync) – This instruction provides ordering for all effects of all instructions executed by the processor. It is used to synchronize the instruction context, such as memory translation, endianness, cache coherency, etc. Instruction pipelines are flushed when an isync is performed, and the next instruction is fetched in the new context. This instruction is useful for self-modifying code.

Memory Management

Memory management is used to translate logical (effective) addresses to physical (real) addresses. Memory management units (MMUs) are also used to control storage attributes, such as cacheability, cache write-though/write-back mode, memory coherency, and guardedness. There are two primary approaches; one defined by PowerPC classic in the 600/700 family of processors and another used by the 400 family and Book E specification. In both cases, the architecture defines a unified MMU, which has traditionally been implemented as independent instruction and data MMUs, enabled via the MSR [IR,DR] bits, respectively. Below is an overview of the two approaches.

PowerPC Classic MMU

The PowerPC Classic MMU was designed primarily for demand page operating systems such as UNIX or MacOS. There are two translation mechanisms, one for block address translation, and another for page tables. Block address translation is performed using eight pairs (upper and lower) of address translation registers, four for instruction addresses (IBATU/L 0-3), and four for data accesses (DBATU/L 0-3). The BAT registers define page sizes ranging from 128KB to 16MB.

For systems requiring more translations than are found in the allocated BAT registers, page table translation is provided. A 32-bit effective address is translated to a 52-bit virtual address, and is then translated into a physical address. One of 16 segment registers (SR0-

SR15) provide virtual address and protection information. Page Table Entries (PTEs) provided physical address and page protection information. The architecture allows for implementations to provide translation-lookaside buffers (TLBs) to speed the translation process, but does not define them. The page-tables are typically programmed by the operating system and their discussion is beyond the scope of this article.

400 Family/Book E MMU

The Book E carries on the idea of a flexible MMU structure for embedded systems. Page sizes are programmable; a page can be large (up to a terabyte in the Book E architecture) to simplify software and minimize the number of entries, or as small as 1KB, to avoid wasting memory space. In addition to normal protection and translation mechanisms, endianness is defined by a page attribute. TLB misses result in an exception; it is under software control to handle the page miss algorithm. A TLB search instruction, **tlbsx**, assists in searching the entire TLB array in a single cycle.

Interrupts

The PowerPC architecture provides a minimal hardware scheme for saving state on interrupts. The only registers that are saved are the IAR and MSR. Interrupt enable bits are disabled for the interrupt type that occurred in order to prevent a second interrupt from occurring before saving the context. Software must save all necessary registers – these typically include all user-mode registers and possibly certain supervisor mode SPRs. Exception-state saving is typically performed by an operating system, but note that for small exception vectors, time can be saved by only saving registers that would otherwise be corrupted. Operating systems must take a more universal approach and save all registers that may be necessary, even if some wind up not being touched by a particular exception handler.

PowerPC Classic Exception Vector Processing

A single interrupt hierarchy is defined. When an interrupt occurs, Save/Restore Register 0 (SRR0) is loaded with the address of where processing should resume after the exception, and the machine state register is saved to SRR1. SRR0 may be loaded with the current IAR or in some cases the next instruction. Interrupt vectors are located at either a high address (0xFFFn_nnnn if MSR[IP=1] or low address (0x000n_nnnn if MSR[IP=0]), depending on the instruction prefix bit in the MSR. The interrupt type determines the lower 5 bits of the vector. When processing is completed, an **rfi** instruction is executed to restore the IAR and MSR to the saved values in SRR0 and SRR1.

400 Family and Book E Exception Vector Processing

Both the IBM 400 family and Book E define a two-level interrupt hierarchy: a non-critical interrupt class, and a critical interrupt class. The non-critical class registers work as previously described for PowerPC classic. For critical interrupts, the IAR and MSR are saved to separate registers (SRR2 & SRR3, respectively for the 400 family, and CSSR0 & CSSR1 for Book E). When a critical exception is completed, an rfci instruction is executed to properly restore the machine. By having a dual-level interrupt scheme, non-critical interrupts can be more easily debugged. More than two sets of interrupt vectors are possible – for the 400 family, the upper 16 bits of the exception vector is contained in the Exception Vector Prefix Register (EVPR). For Book E, all 16 exceptions can have the upper half of the exception vector mapped to a different location through the use of 16 Interrupt Vector Prefix Registers (IVPR0-15).

Stack

The PowerPC architecture has no notion of a stack for local storage. There are no push or pop instructions and no dedicated stack pointer register defined by the architecture. However, there is a software standard used for C/C++ programs called the Embedded Application Binary Interface (EABI) which defines register and memory conventions for a stack. The EABI reserves GPR1 for a stack pointer, GPR3-GPR7 for function argument passing and GPR3 for function return values. Assembly language programs wishing to interface to C/C++ code must follow the same standards to preserve the conventions.

Caches

The PowerPC architecture contains cache management instructions for both user-level and supervisor-level cache accesses. Cache management instructions are found in Table 4 below.

Instruction	Mode	Implementation	Function
dcbf	User	All	Flush Data Cache Line
dcbi	Supervisor	All	Invalidate Data Cache Line
dcbst	User	All	Store Data Cache Line
dcbt	User	All	Touch Data Cache Line (for load)
dcbtst	User	All	Touch Data Cache Line (for store)
dcbz	User	All	Zero Data Cache Line
dccci	Supervisor	IBM 4xx	Data Cache Congruence Class Invalidate
icbi	User	All	Invalidate Instruction Cache Line
icbt	User	4xx / Book E	Touch Instruction Cache Line
iccci	Supervisor	IBM 4xx	Instruction Cache Congruence Class Invalidate

Table 4 - Cache management instructions

Care should be taken when porting cache manipulation code to a different PowerPC implementation. Although cache instructions may be common across different implementations, cache organization and size may likely change. For example, code that makes assumptions about the cache size to perform a flush may need to be modified for other cache sizes. Also, cache initialization may vary between implementations. Some provide hardware to automatically clear cache tags, while others require software looping to invalidate cache tags.

Self-Modifying Code

While it is not a recommended practice to write self-modifying code, sometimes it is absolutely necessary. The following sequence shows the instructions used to perform a code modification:

1. Store modified instruction.
2. Issue **dcbst** instruction to force new instruction to main store.
3. Issue **sync** instruction to ensure DCBST is completed.
4. Issue **icbi** instruction to invalidate instruction cache line.
5. Issue **isync** instruction to clear instruction pipeline.
6. It is now OK to execute the modified instruction.

Timers

Most implementations have provided a 64-bit timebase that is readable via two 32-bit registers. The amount the timer increments varies across families, as well as the SPR numbers and instructions to access the timebase. Therefore, care should be taken when porting timer code across implementations. Additional timers may also vary, but most provide at least one kind of decrementing programmable timer.

Book E Timers

Both the IBM 400 family and Book E define the following timers in addition to the timebase: a 32-bit programmable decrementer (DEC in Book E, PIT for the 400 family) with an auto-reload capability, a fixed-interval timer (FIT), and a watchdog timer (WDT) for system hang conditions.

Debug Facilities

Debug facilities vary greatly between implementations. Original PowerPC 600 family parts had only one instruction address breakpoint. PowerPC 700 family parts have added a single data address breakpoint. PowerPC 400 family parts have much more robust debug capabilities, including multiple instruction address breakpoints, data address breakpoints, and data value compares. Other features may include breakpoint sequencing, counters, ranges, and trace capabilities.

Maintaining Code Compatibility

PowerPC users who expect to program for more than one implementation typically ask for tips on maintaining code compatibility. The following are some suggestions to help minimize porting problems:

- Use C code whenever possible. Today's C compilers can produce code that is comparable in performance to hand-assembly coding in many cases. C code, being Book I code, will guarantee code portability.
- Also, try not to embed processor-specific assembly instructions in C, as they'll be harder to find. Separate processor-specific code that is known to contain device dependent registers or instructions. These are typically things like boot up sequences and device drivers, but also may include floating point code (including long long types). Keep them well documented as to assumptions and dependencies.
- Use the PVR, but only when appropriate. Common code across minor variations of implementations is good, and the PVR can be used for decision making. But in the case where major modifications are necessary (for example, 7xx versus 4xx MMU code), separate code bases are recommended.

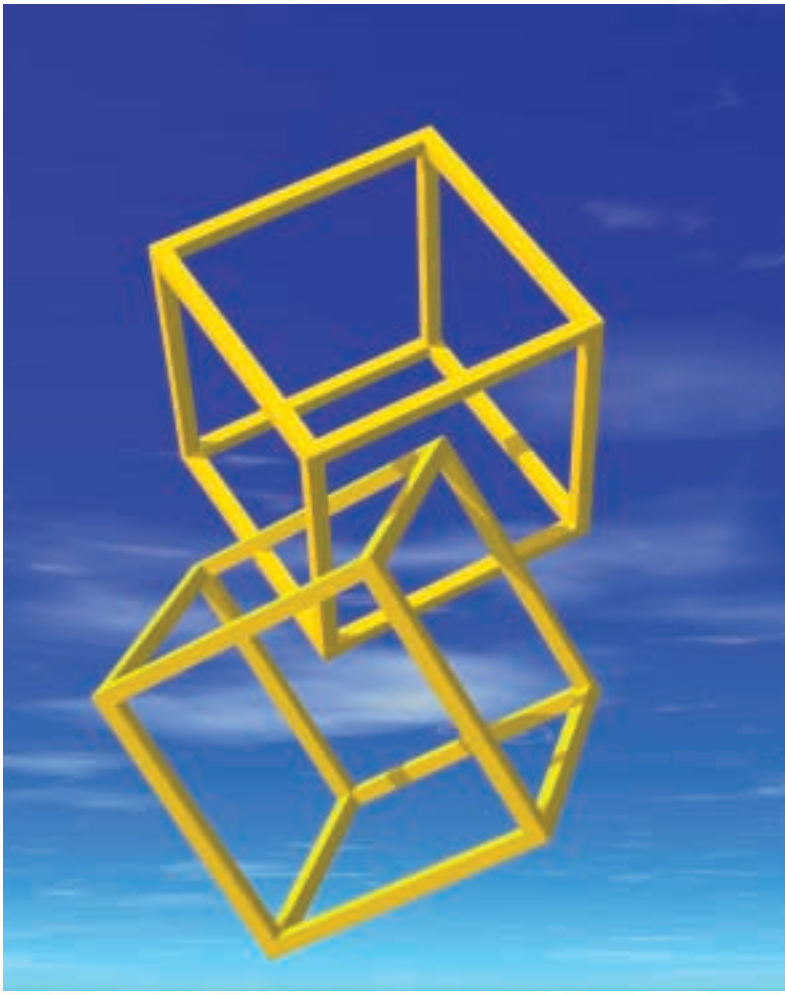

Summary

This completes an introduction to the PowerPC programming model. IBM hopes you have found this of value, and that it adds to the success of your development programs. For more information, go to: *www-3.ibm.com/chips/products/powerpc/newsletter/*

IBM's New PowerPC Strategy Roadmap

It's exactly what you'd expect from IBM.

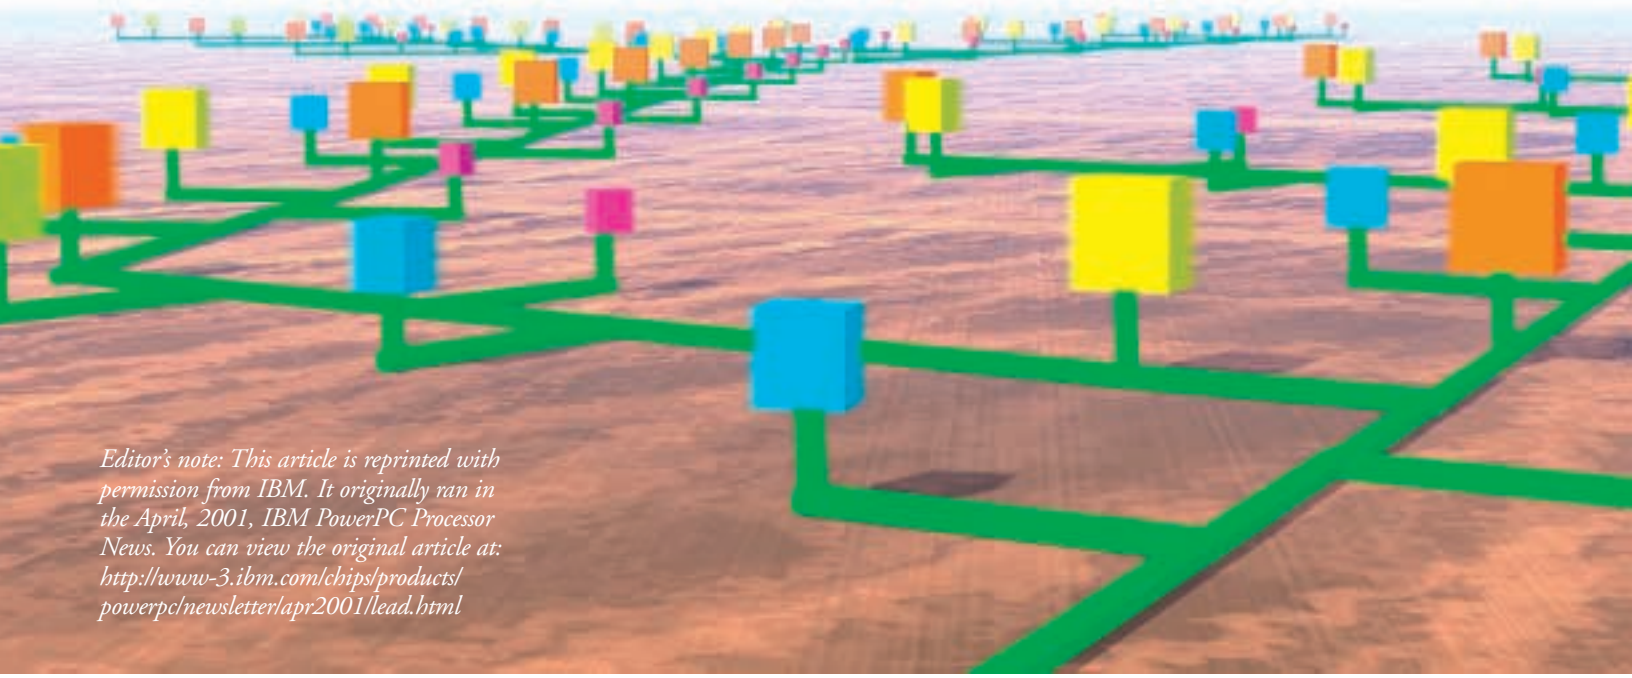

Editor's note: This article is reprinted with permission from IBM. It originally ran in the April, 2001, IBM PowerPC Processor News. You can view the original article at: http://www-3.ibm.com/chips/products/powerpc/newsletter/apr2001/lead.html

by Kalpesh Gala and Mike Vowell
IBM

It is not by accident that ideal targets for IBM PowerPC™ technology include wired and wireless networking, storage, and pervasive computing applications. It's all part of a master strategy to focus both today's and tomorrow's industry leading technology on PowerPC products – to meet the performance, power, and price needs of these ever-dynamic and diverse applications.

The focus of IBM's strategy centers upon both technology and design expertise. Underlying this focus is a strong and long-term commitment to leveraging the PowerPC architecture and designing with standard interfaces, such as RapidIO™, PCI/X, and Ethernet. The goal of this strategy is to assure the continuation of a product portfolio heritage of being core-based, power-efficient, scaleable, and software transparent.

The melding of IBM's advanced technology, PowerPC products, and SoC capability will establish a hallmark in the battle for mindshare and marketshare. As market opportunities evolve for higher performance and/or lower power devices, IBM will be poised to address these seemingly insatiable needs.

THE RAPIDIO INTERCONNECT ARCHITECTURE, DESIGNED TO BE COMPATIBLE WITH MOST POPULAR INTEGRATED COMMUNICATION PROCESSORS, HOST PROCESSORS, AND NETWORKING DIGITAL SIGNAL PROCESSORS, IS A HIGH PERFORMANCE, PACKET-SWITCHED, INTERCONNECT TECHNOLOGY. IT ADDRESSES THE HIGH-PERFORMANCE EMBEDDED INDUSTRY'S NEED FOR RELIABILITY, INCREASED BANDWIDTH, AND FASTER BUS SPEEDS IN AN INTRA-SYSTEM INTERCONNECT. THE RAPIDIO INTERCONNECT ALLOWS CHIP-TO-CHIP AND BOARD-TO-BOARD COMMUNICATIONS AT PERFORMANCE LEVELS SCALING TO TEN GIGABITS PER SECOND AND BEYOND.

IBM's PowerPC cores, microprocessors, and integrated products meet the unique needs of an increasingly diverse marketplace. PowerPC chips offer state of the art

technology in a variety of configurations to provide the optimal mix of performance, power, functionality, and size. And because of our core-based design philosophy, IBM customers can quickly and easily differentiate their products in the marketplace, and still maintain flexibility and software transparency across generations of devices.

IBM has been a leader in awarded patents for many years, largely as a result of IBM Microelectronics contributions. Among the more recent of these contributions are several dozen patents directly related to three chip breakthroughs – silicon germanium (SiGe), silicin-on-insulator (SOI), and Low-K Dielectric.

Just as IBM led the industry with its copper process technology, these new process technologies will be the catalyst for even greater performance and low-power advances in the semiconductor industry.

Key to enabling a high level of flexibility is the IBM CoreConnect™ on-chip bus architecture, which is becoming a defacto industry standard. CoreConnect provides a standardized method for assembling pieces of chip designs from diverse suppliers to facilitate an open SoC design process that encourages the development of reusable IP. Currently, this bus structure is licensed by over 40 IP providers, and is the basis for numerous IBM standard, application specific, and custom devices. Forthcoming enhancements include higher bandwidth and crossbar functionality, which will marry well with our next-generation CPU cores and our planned addition of performance-enhancing IP like RapidIO and high-speed serial ports.

Enhancing IBM's strategy is a commitment to the PowerPC architecture. The latest enhanced version, called PowerPC Book E, has been refined to provide 64-bit capabilities, increase flexibility, and address the unique demands posed by embedded systems. Book E is a new definition of the PowerPC architecture, one that maintains compatibility with applications developed for the original PowerPC architecture.

What's technology got to do with it?

Almost everything. It's what boosts performance... throttles power consumption... and enables smaller devices. In a nutshell, technology is a key ingredient to the ever-elusive faster, smaller, and cheaper semiconductor solution.

Today's copper process, silicon-on-insulator, and Low-K Dielectric technologies have fostered higher levels of performance and lower power dissipation than were previously attainable. This elevation of performance is currently being realized in IBM's highest performance servers and workstations for data processing and electronic commerce through IBM's stand-alone and system-on-chip (SoC) processors.

The application of new process technologies is crucial to the development of processor engines that will meet the performance and power needs of next-generation applications. And nowhere will the infusion of new technology be more appreciated than in the relatively new "pervasive computing." This includes handheld and embedded products such as smart phones and Internet appliances for business professionals and consumers. In this arena, both "high performance" and "low power consumption" are of utmost importance. To continue meeting this challenge, IBM is focusing its technology resources on PowerPC embedded processors that will distinguish themselves by their enviable power/performance ratios as well as their high integration of critical IP.

Conclusion

As a chief architect of the PowerPC Architecture, IBM has been at the forefront in the evolution of microprocessor design, semiconductor technology, and SoC advances in the industry. IBM likes its position as a pioneer and industry leader – and has no desire to relinquish its position. You are invited to meet IBM at any of its milestones, as it advances toward an ultrascalar, 2+ GHz-processor engine. It's exactly what you would expect from IBM.

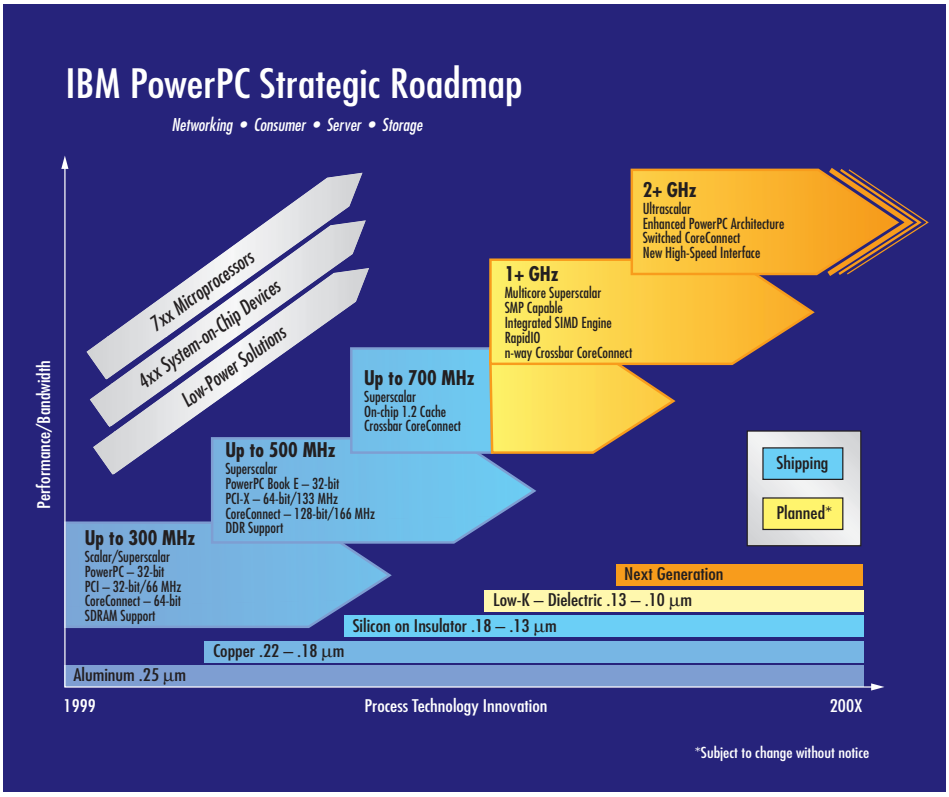

Get up to Multi-Gigabit Speed with the SPECCTRAQuest Design Kit

Learn how to implement Rocket I/O multi-gigabit serial transceivers in the new Virtex-II Pro Platform FPGA.

by Donald Telian
Technologist
Cadence Design Systems
donaldt@cadence.com

Mark Alexander
Product Application Engineer
Xilinx, Inc.
mark.alexander@xilinx.com

With the introduction of the new Virtex-II Pro™ FPGAs, high-speed Rocket I/O™ serial transceivers are ready to find their way into hundreds of new applications. Are you ready? What used to be confined to a few exotic chips, laden with pages and pages of design and implementation guidelines, is now available to everyone. Xilinx has made it possible to integrate multi-gigabit transceivers (MGTs) into your FPGA, but how will you integrate that high-speed FPGA into your system of printed circuit boards (PCBs)?

In an effort to avoid multi-gigabit headaches, Xilinx and Cadence Design Systems have assembled SPECCTRAQuest™ MGT Design Kits to help you implement the new Rocket I/O MGTs effectively in your system. Whether you need to craft a custom MGT interface and develop your own PCB/backplane/cabling guidelines, or you simply need to apply your MGTs in a standard configuration, the kit gets you moving towards a solution within minutes.

SPECCTRAQuest MGT Design Kit

Pre-configured circuits in the design kits are ready to simulate for both typical MGT chip-to-chip and backplane PCB interfaces. There's no time wasted hunting for models, testing and correlating them, or figuring out how to connect them together. It has all been done for you. And because the simulation environment is graphical, adapting the circuits for your unique application is as simple as dragging and dropping.

Better yet, the models of the active Rocket I/O MGT circuitry are transistor-level silicon models that have been correlated by Xilinx to match both the actual silicon design and empirical data. This ensures that your system implementation is designed with the most advanced and accurate models available. Add to that fully coupled frequency-dependent lossy package, PCB trace, and connector models, and you're ready to carefully characterize the signal integrity and degradation issues that are inherent in this type of design.

Figure 1 shows a pre-configured simulation drawing for chip-to-chip applications. Although this is actually a 500+ node simulation, it has been simplified through the use of subcircuits and black box models so it can be more easily modified. Just point and click to change trace parameters, physical connections, or other aspects of the circuit.

When your simulations are done and it's time to layout your PCB, the MGT kit has sample footprints and constraint files to ensure a successful layout. It is essential to bind all high-speed constraints into your

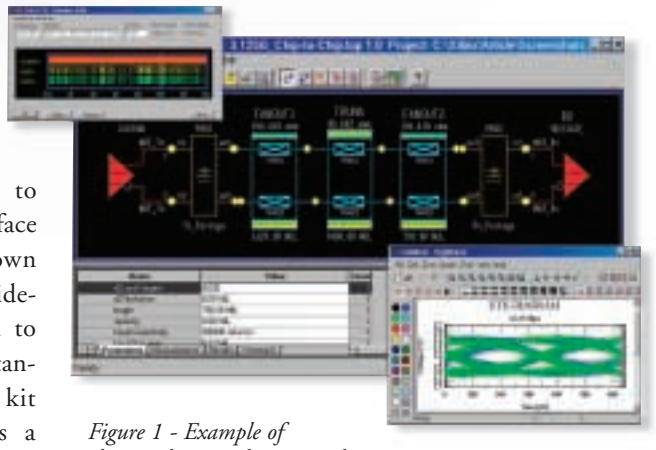

Figure 1 - Example of chip-to-chip simulation topology

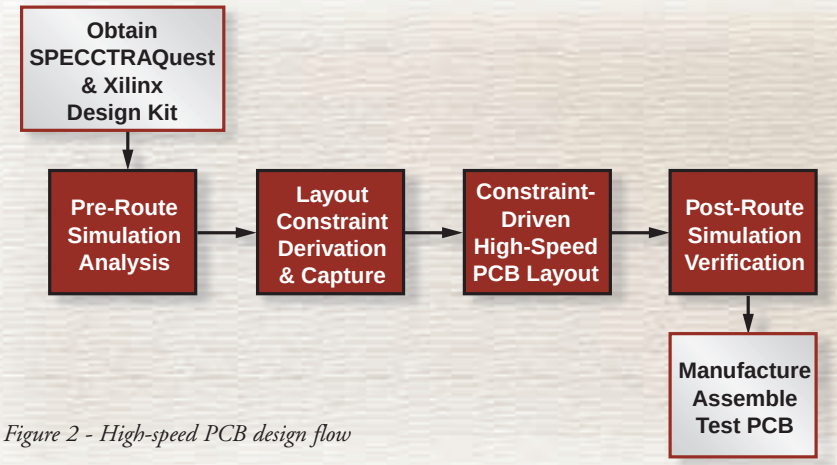

Figure 2 - High-speed PCB design flow

PCB database to automate and properly constrain the layout process. If you don't bind your constraints, your layout, post-layout simulation, and actual system behavior will not be as clean as you would want.

To help you along the learning curve associated with designing multi-gigabit links, the kits also contain tutorial "movies" you can run on your PC to clearly illustrate how to use the various aspects of the kit throughout your design process.

SPECCTRAQuest's Features

SPECCTRAQuest is an integrated design tool that allows you to include both your MGT silicon models and PCB databases in one simulation. You can easily set up extensive "sweep" simulations to sweep circuit variables through a range of values to test out different implementation options.

Simulation of trace parameter variations (such as loss tangents, skin effect, dielectric constant, and other variations), crosstalk, power noise, deterministic/random jitter, and other effects are all included.

Stimulate your circuit with pre-coded 8b/10b pseudo-random bit sequence (PRBS) patterns to create inter-symbol interference (ISI) and study its effect on signal integrity. With the push of a button, you can plot the resulting waveform as an eye diagram so you can quickly quantify the signal degradation from transmitter to receiver. SPECCTRAQuest also provides automated measurements of each circuit's performance. These measurements can be sorted so you can quickly select the best configuration.

After you have determined the best implementation configuration, you can use Cadence's powerful constraint management (CM) tool to electronically capture and manage your layout constraints. The CM tool interacts with the other Connected Team Design Tools (including Concept HDL, Allegro, SPECCTRA™, and SPECCTRAQuest) to enable your whole design team to work concurrently and share data.

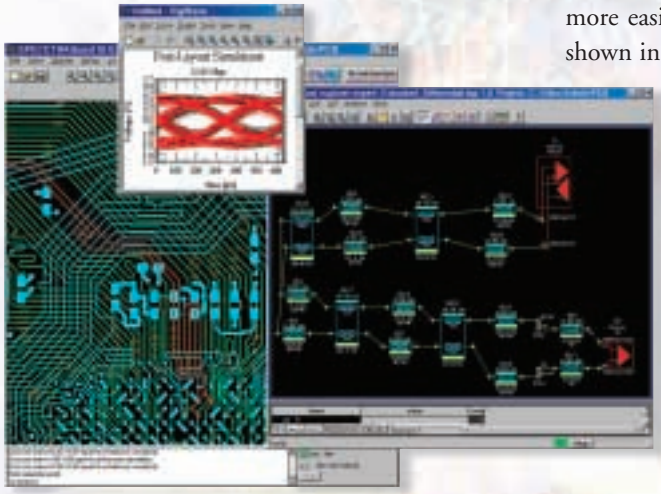

Figure 3 - Simulation from layout and extracting a differential net

SPECCTRAQuest has a unique way of superimposing advanced simulation on top of high-speed PCB layout. For more information on the features available in SPECCTRAQuest and other Cadence PCB tools, please refer to *www.specctraquest.com* and *pcb.cadence.com*.

MGT Design & Implementation Flow

With both the SPECCTRAQuest software and MGT design kit in place, you're ready to exercise the advanced, high-speed design flow shown in Figure 2 to ensure your MGTs are implemented correctly.

The flow begins by configuring a simulation of your unique application. Use the pre-configured topologies in the design kit as a starting point. Simulate your application pre-route by sweeping through all the implementation options and manufacturing tolerances to ensure that the multi-gigabit signals are transmitted and received within tolerances. Use the test patterns and post-processing tools in the kit to help with this phase.

Once an acceptable solution is found, capture the layout constraints in electronic topologies that can be bound into the layout process. Once again, you can use the kit examples to learn how. Perform constraint-driven layout, and then do post-layout simulation on the routed nets to verify that signal transmission is still within tolerance. If the layout constraints were applied and followed, this will be a simple verification step. If a problem is found, the differential net can be extracted to an electrical view to more easily identify the cause of failure, as shown in Figure 3.

Conclusion

Although using MGTs in your high-speed design can bring higher performance to your FPGA application, you must be careful in performing the implementation. By correctly using the SPECCTRAQuest MGT Design Kit assembled for this task, you can avoid common problems and shorten your design cycle. The kit includes layout guidelines, constraints, topologies, scripts, and utilities that will help you integrate Rocket I/O MGTs into your high-speed Virtex-II Pro designs.

The MGT design kit is available for free download from the Spice Suite at www.xilinx.com. Registration is required if you are not already a Xilinx customer

Best-of-Class Embedded Software Development Tools

A compiler tuned for a specific processor architecture automatically produces tight, fast code.

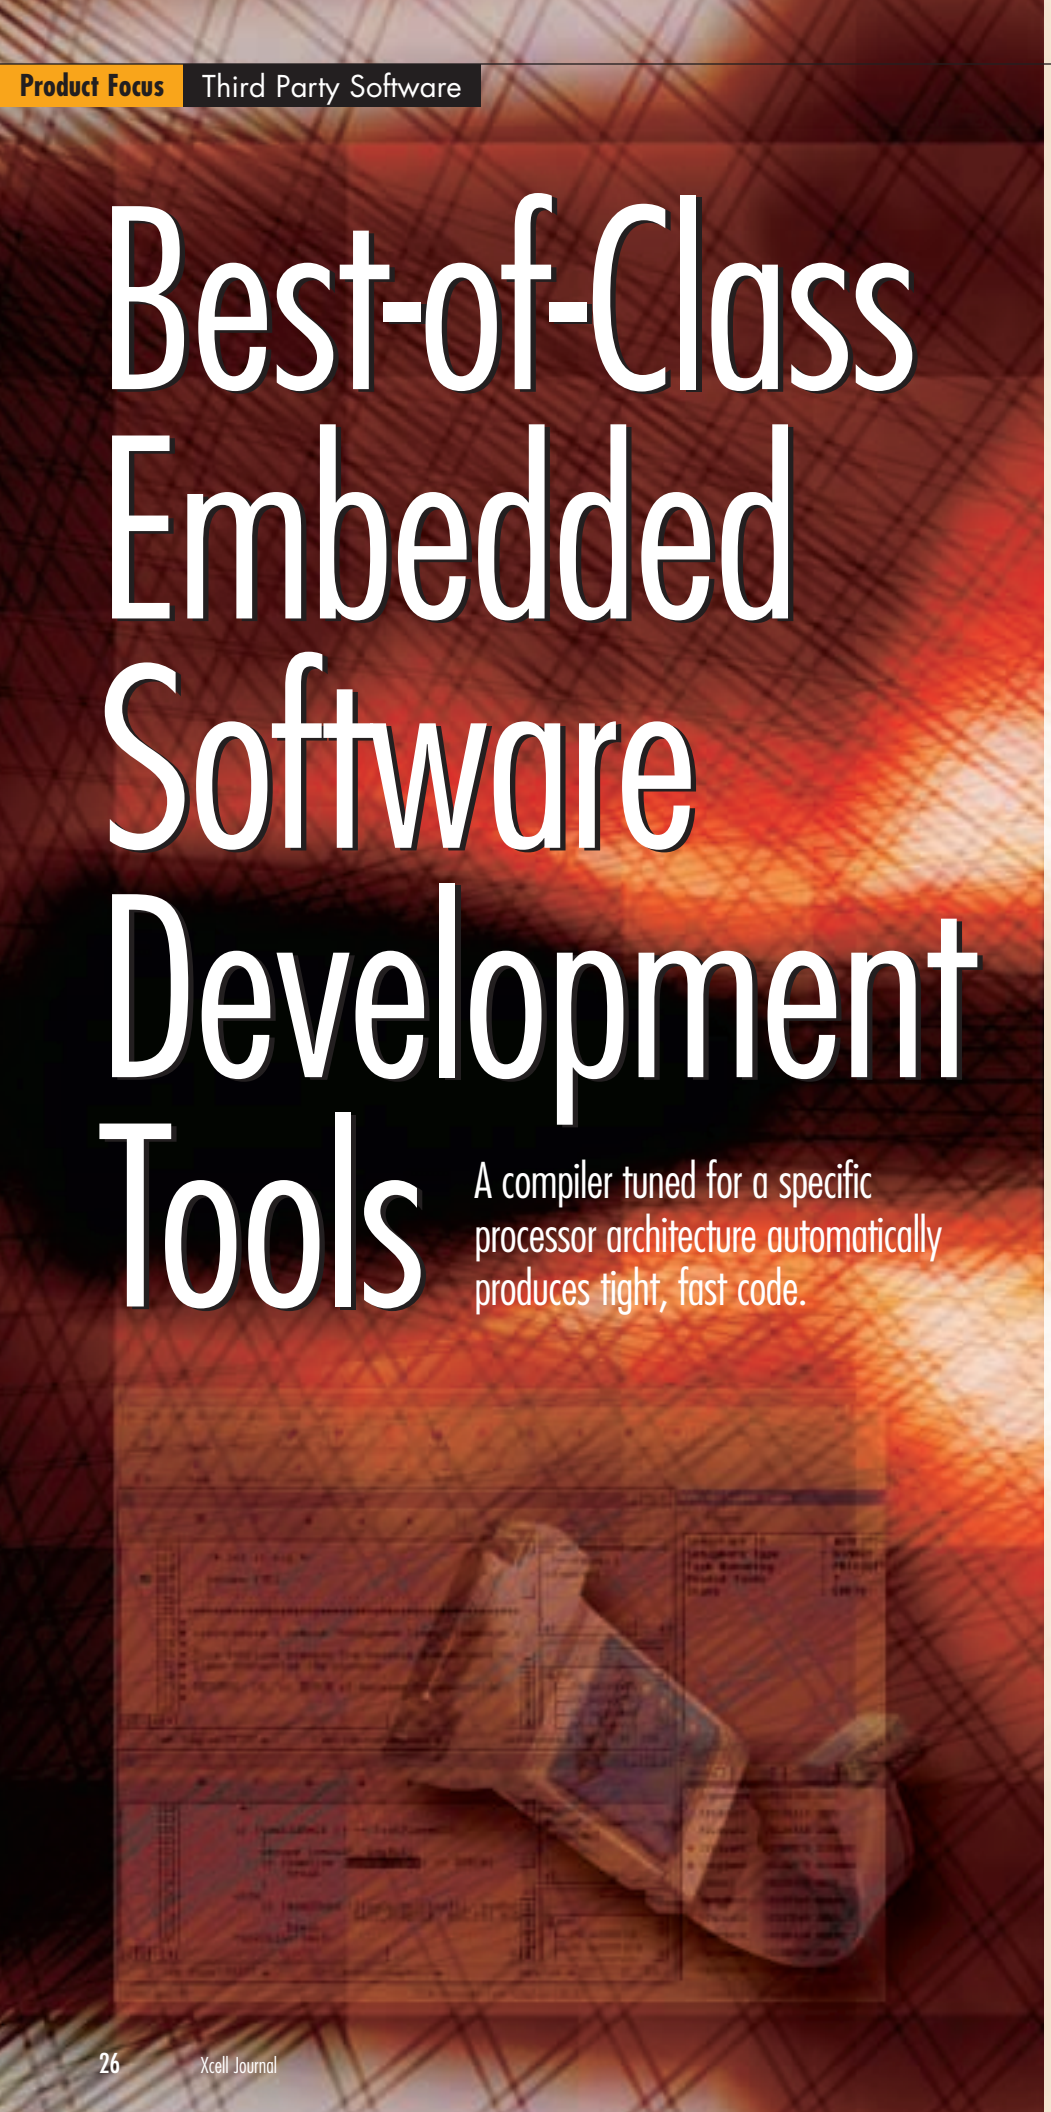

by Jay Gould
Embedded Partners Program Manager
Xilinx, Inc.
jay.gould@xilinx.com

Xilinx Platform FPGAs introduce a completely flexible and programmable platform for creating custom and upgradeable embedded solutions. However, using high-performance processor cores can be a challenge. To maximize your productivity, you need best-of-class embedded software development tools.

Time-to-market pressures and product differentiation goals often create conflicts with engineering schedules. That's why more companies are moving away from home-grown tool environments and their associated support problems. In this competitive climate, it's usually better for you to use the best, commercially available tools – and focus your engineering resources on unique, value-added product design.

Virtex-II Pro Requires Embedded Software Tools

The term "embedded software tools" most often applies to the tools required to create, edit, compile, link, load, and debug high-level language code (usually C/C++) for execution on a processor engine. The processor could be hard or soft; 8-, 16-, 32-, or 64-bit; high or low performance, and so on, but the basic development flow is generally the same. With Virtex-II Pro™, you can target design modules for either silicon hardware in FPGA logic gates, or as software applications run on processor engines like the embedded IBM PPC405 hard core. Since hardware engineers, software engineers, firmware engineers, system architects, and others may all target the Virtex-II Pro, Xilinx has a "market leader" tools strategy to appeal to these different camps. This strategy also has the added advantage of appealing to the largest installed base of embedded users.

Xilinx could have created new embedded software tools from scratch, (in fact, we employ more software engineers than most "software" companies), but instead Xilinx chose to launch the Virtex-II Pro with "Xilinx versions" of established third-party tools. Therefore you don't have to embrace completely new development methodologies and you can port existing designs into the Virtex-II Pro fabric. With tens of thousands of engineers already using these tools, using a complementary model is far more constructive than creating a new technology.

For an embedded processor core to run machine code, the algorithm must be entered in an HLL (high level language) such as C or C++ rather than in an HDL (hardware description language). A compiler is used to translate that HLL design into specific (PPC405) binary code that can be executed on that particular processor.

Bad compilers produce poorly optimized, bloated code that takes up a lot of memory and runs slowly. A best-of-class compiler is one which has been tuned for a specific processor architecture and automatically produces tight, fast code. This saves memory for the design and spares you days of hand-optimizing, trying to manually accelerate your code. For the integration and debugging stages, a software debugger is run on a host computer or workstation, and controls the program execution of the embedded target over a hardware connection probe. With a Virtex-II Pro and the PPC405, this communicates directly with the processor core using a JTAG port.

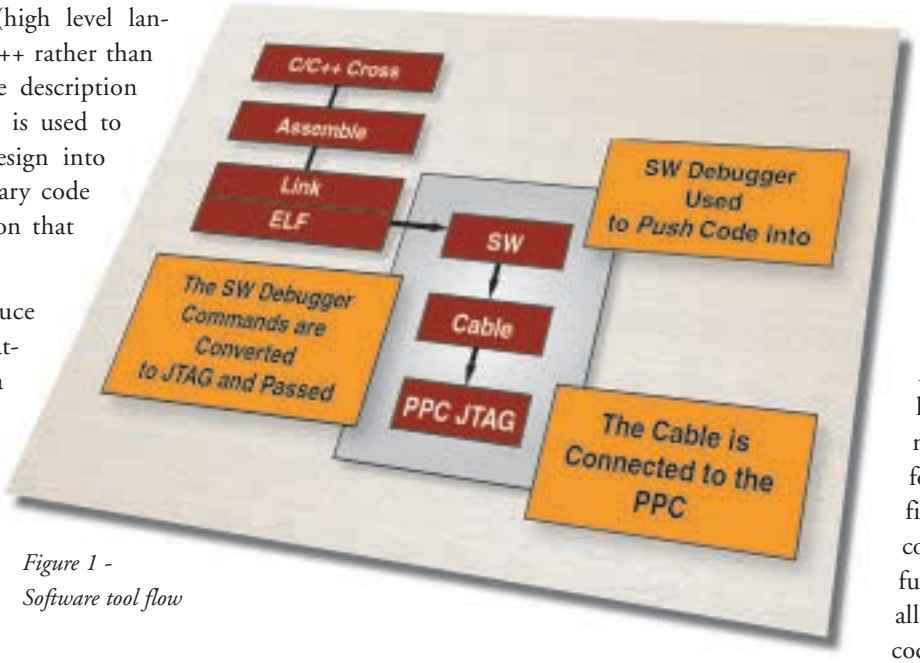

Figure 1 - Software tool flow

Partnering with Wind River Systems

Xilinx worked with Wind River Systems, the market leader in the embedded software industry, to provide a robust set of software tools for targeting the PPC405 in Virtex-II Pro FPGAs. Wind River Systems has a commanding market share with their broad range of development tools, Real Time Operating Systems (RTOSs), and middleware solutions. A specific Virtex-II Pro version of the Wind River tools (compiler, software debugger, and JTAG run control hardware probe) have been created for Xilinx distribution via an OEM agreement.

The Wind River tools are optimized for the PPC405. The Virtex-II Pro versions of these tools are created with some basic design size limitations so that Xilinx can package a lower-price entry point tool suite through the Xilinx sales channel. You can buy the unrestricted versions (or upgrades) from the Wind River sales channel. Wind River already supports the PPC405 and will also offer numerous other embedded tools.

Wind River makes numerous tools themselves and have a huge partnership program where third parties build other tools to integrate within the Wind River Tornado IDE (integrated development environment). This opens the market to numerous other software tool suppliers who can specialize in niche areas, providing extra value to end users. Wind River owns multiple Real Time Operating Systems (RTOSs), but their main market focus is on an RTOS product called VxWorks, which is also the market leader.

The Wind River OEM tools for Virtex-II Pro and PPC405 include:

- **Diab XE** – compiler
- **SingleStep XE** – software debugger
- **visionPROBE II XE** – JTAG run control hardware connection probe

"XE" stands for "Xilinx Edition".

Optimized Compilers Provide High-Performance Code

The Diab C/C++ Compiler consistently scores high in the annual market research results for best compilers. The fact that Diab is highly optimized for the PowerPC and has long been a market leader made it a perfect fit for the Virtex-II Pro and embedded PPC405 core.

A compiler turns high level language into machine executable code for a processor. Providing fine grained control of the compiler options, Diab is a full featured product which allows you to balance speed, code size, and memory usage for your applications. At the front end of the Diab product is a language parser which creates a language-independent representation of your code, and this unleashes the power of the five different back-end optimizer stages of the compiler. Diab provides optimization for global, code selector, code generator, peephole, and instruction scheduler stages as well as supporting architecture-specific features such as the SPRs/DCRs on the IBM 405 processor.

Both the Diab Compiler and the SingleStep Software Debugger are considered "best of class" products in the PPC embedded domain. The Xilinx PowerPC OEM software product is based on these top-of-the-line products.

Debugging Embedded Software on the Real Target Hardware

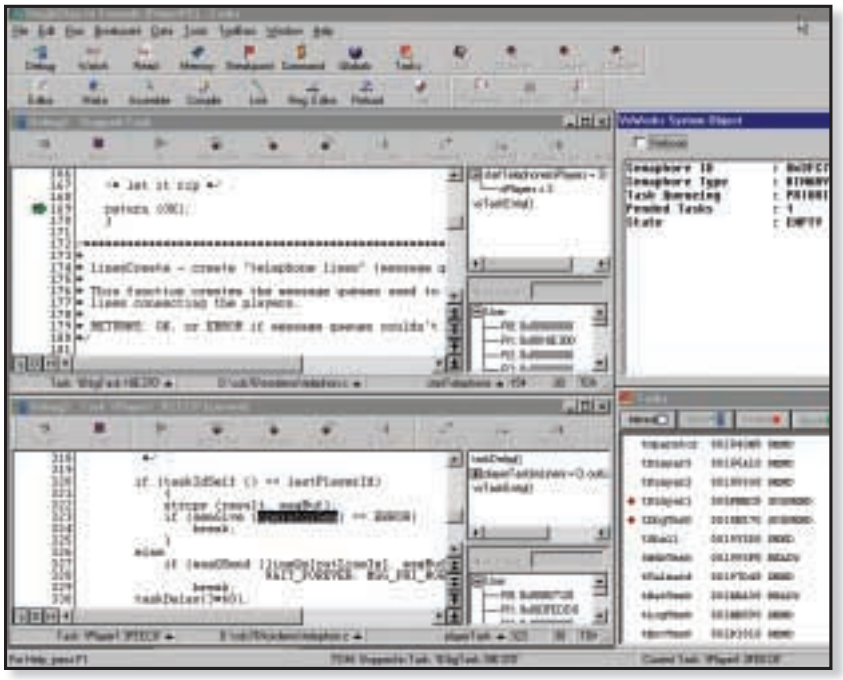

Figure 2 - SingleStep software debugger

SingleStep is a multi-windowed and full featured embedded software debugging tool that is far superior to command line tools. This product is ideally suited for board/hardware development, driver/firmware development, and software application debugging (thus reducing the requirements from two or more possible tools to just one). Software debuggers need to provide basic "run, start, stop" control to debug an embedded system, but SingleStep exceeds these expectations as one of the most feature rich embedded debuggers available.

SingleStep supports real-time target control, high-speed downloads, built-in hardware diagnostics and Flash memory programming. It provides a unique processor-specific register interface to enable configuring and initializing integrated peripherals, and a command line interface with scripting language to automate testing.

This debugger provides superior PPC405 support through hardware breakpoints – including four hardware instruction address, two data instruction address and two data value breakpoints – and a register window. In addition to on-chip trace, SingleStep provides visibility for instruction completion, branch taken, exception taken, data address compare, and other debug events.

Connect Software Tools to the Embedded Hardware

To download new FPGA images or run the host debug software tools on the real embedded system, visionPROBE II provides a high-speed parallel connection between the SingleStep debugger on the host and the Virtex-II Pro target. This product allows JTAG run control of the target for system debugging and, via hardware caching logic, the visionPROBE II can execute high-speed downloads up to 400 KBps.

The probe facilitates register initialization, hardware diagnostics, Flash programming, and hardware breakpoints. Additionally, this one probe allows for FPGA image downloads and ChipScope™ operation on the same target without having to add or change cables.

Figure 3 - visionPROBE II

RTOS Validation

Xilinx will provide some validation support for Real Time Operating System (RTOS) environments. Depending on the complexity of your application, an embedded system design may or may not require a hard-real-time operating system. If you are experimenting with a gate-consuming protocol or algorithm, some simple C/C++ code targeted for a Xilinx MicroBlaze™ soft processor may be all that is required. For other more complex applications, which have tight requirements for fast interrupts, maximum uptime and minimum latency, the PPC405 and a robust RTOS may be required.

Wind River – VxWorks

Due to the close partnership with Wind River, Xilinx will validate that the Virtex-II Pro development boards, and our tools, will work seamlessly with the Wind River VxWorks RTOS. In fact, the entire support of the Virtex-II Pro will be validated in the Wind River environment, including the certification of a Xilinx "Board Support Package" (BSP) that matches with our reference board. The certification is a formal process that is executed by Wind River and through which they publicly acknowledge a working configuration of our tools, boards, and devices.

Conclusion

If full programmable embedded system design with the Virtex-II Pro is what you are after, adding the software development flow to your FPGA logic tools is the next step. Rather than create a new methodology and introduce a new tool suite to a mature market, Xilinx has chosen to partner with the market leader to access the best-of-class tools for the PowerPC. Proven and established tools allow you to focus on adding value to your product design, without learning the nuances of immature or un-optimized tools. Read more about the Xilinx embedded software solutions at: *www.xilinx.com/processor*

Three powerful reasons to make Wind River your embedded solutions partner.

Discover why Wind River's unmatched experience, innovation, and vision provide the ideal roadmap for success.

For nearly twenty years, Wind River's embedded expertise has provided direction for advanced technology companies worldwide. Today, this experience has made us a leader in a variety of growing markets, offering exceptional performance, selection, and reliability.

- End-to-end development solutions
- Excellent time to market
- Strong joint development programs
- Unequaled customer service and support
- Industry-leading technology
- Complete life-cycle support and upgradability

Let us show you the fastest route to superior embedded performance in your application. Visit us online for more information at www.windriver.com

See us at the Xilinx Programmable World 2002

WindRiver®

NETWORKING | AUTOMOTIVE | CONSUMER ELECTRONICS | INDUSTRIAL | AEROSPACE/DEFENSE | SERVICES

ISE 4.2i Expands High-Powered Design Productivity

ISE 4.2i extends its coverage to new Xilinx logic devices and even greater design tool productivity.

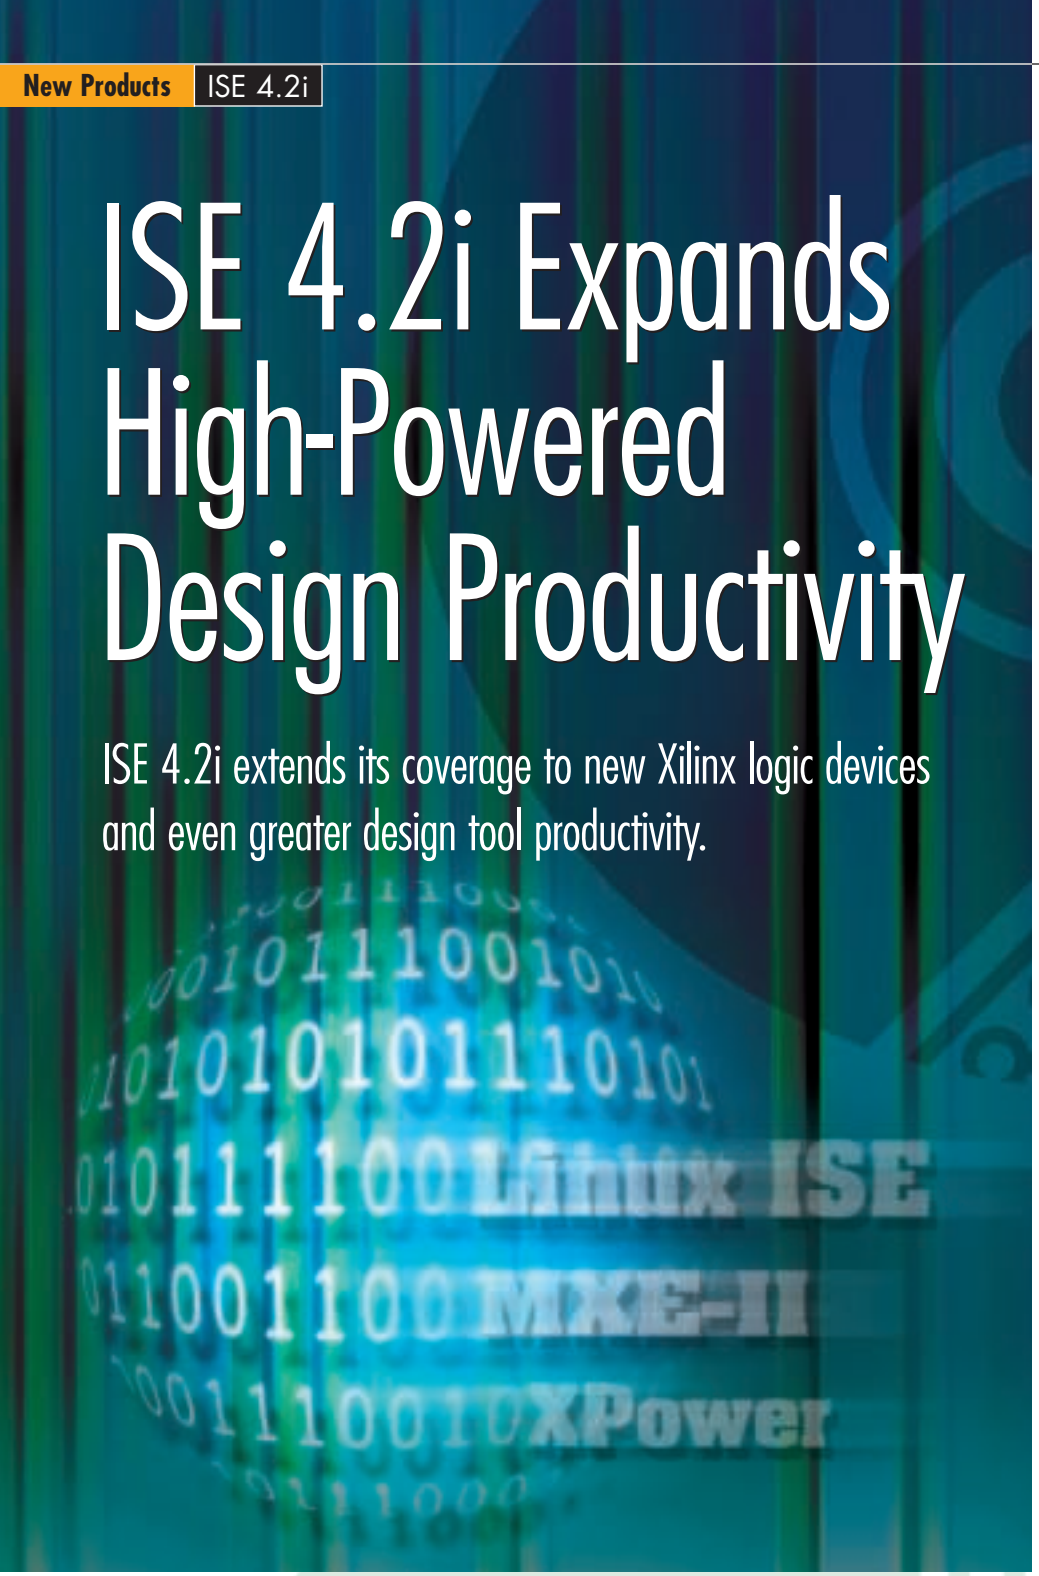

by Lee Hansen
Software Product Marketing Manager
Xilinx, Inc.
lee.hansen@xilinx.com

In February, Xilinx released Integrated Software Environment version 4.2i logic design software. ISE 4.2i delivers several new capabilities that make it the most complete package available for programmable logic design. Building on the advances of ISE 4.1i, released this past August, version 4.2i offers several advantages in particular to designers looking at embedded logic systems and high-speed signal challenges:

- Linux ISE
- Expanded device support
- High-speed design support
- Partial reconfigurability
- Streamlining XPower performance
- MXE-II

Linux ISE

ISE version 4.2i is the first version of Xilinx design software to run on the Linux operating system. Now, you can run Xilinx implementation tools on Linux Red Hat version 7.2, including the Wine windows application layer. With the addition of the Linux operating system, you now have one of the widest choices possible in design system platforms and implementation tools for programmable logic design.

Expanded Device Support

ISE 4.2i includes the latest device support for all of the Xilinx families, including the new Virtex-II Pro™ product line.

- The Virtex-II Pro Platform FPGA is the only programmable device available with 3.125 Gbps serial I/O – and anywhere from zero to four IBM PowerPC™ 405 microprocessors embedded in the device fabric.
- ISE 4.2i comes ready to implement Virtex-II Pro functions, as well as those available in the recently announced CoolRunner™-II and Spartan™-IIE device families.
- ISE 4.2i also includes the new "-6" speed grade for the industry's fastest Virtex-II FPGAs.
- By delivering FPGA and CPLD device support from a single design product line, ISE gives you an easy way to move up or down device families without having to load or relearn new software.

High-Speed Design Support

With Virtex-II Pro Platform FPGAs, you can now drive design I/O at speeds up to 3.125 Gigabits per second – a first in programmable logic. And ISE 4.2i comes ready to help you realize that high-speed I/O potential quickly and easily. All of the implementation tools, the pin planner, and timing and constraints editors in ISE 4.2i have been enhanced to simplify the implementation of the high-speed I/O of the Virtex-II Pro devices. And the synthesis tools can optimize the paths to and from the multi-gigabit transceivers

(MGTs), a capability unique to ISE 4.2i design software.

The ISE 4.2i libraries also support instantiation of 1, 2, and 4-byte flavors of the high-speed I/O protocols supported by Xilinx.

Xilinx and Cadence Design Systems have jointly developed high-speed design kits. These kits include software components and design aids to help you use the Virtex-II Pro multi-gigabit transceivers. The kits also contain HSpice compatible models for analyzing PCB electrical trace effects coming off the high-speed Virtex-II Pro I/O pins. By combining the new HSpice models for the Virtex-II Pro MGTs, IBIS models for all our FPGA device families and SWIFT models describing the behavior of the PowerPC™ microprocessor and multi-gigabit transceivers, Xilinx provides support for analyzing your FPGA, even after it's been programmed and ready for the board.

Partial Reconfigurability

Continuing its long string of firsts in the programmable logic industry, Xilinx has added a new capability to both the ISE software family, as well as the Virtex-E and Virtex-II product lines – Partial Reconfigurability. Introduced in ISE 4.2i, Partial Reconfigurability allows either a Virtex-II or Virtex-E FPGA to be partitioned so that part of the device can be reprogrammed, while the remainder of the FPGA continues to run.

Partial Reconfigurability works through the Modular Design option for ISE. Using Modular Design, the overall design can be partitioned into sub-modules that can be implemented independently of each other. Once a module is completed, its timing and performance are locked down while the remaining modules are being finished, delivering faster overall design completion.

With the addition of Partial Reconfigurability, the device can also be partitioned where electronic functions make the most sense. Then later on in the field, one partition can be reprogrammed while the remainder of the FPGA continues to run. Products can be updated for new functionality while still in the field, and the product doesn't have to be taken offline while new functionality is loaded.

Streamlining XPower Performance

XPower, an ISE feature that estimates the power consumption of FPGAs and CPLDs, has also been streamlined for better performance. VCD simulation files now read in much faster than previous versions, allowing for quick and accurate setup of input signal activity. And XPower now supports the new CoolRunner-II family of CPLDs.

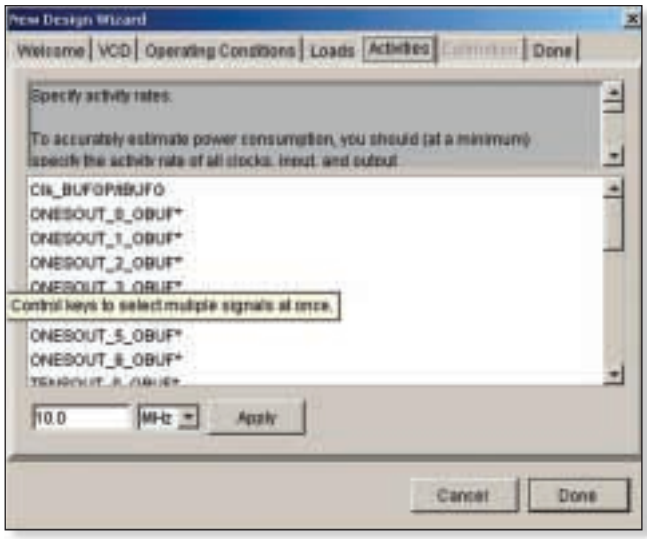

Figure 1 - XPower New Design Wizard

XPower New Design Wizard

Most important, in the ISE 4.2i version of XPower you will also find the "XPower New Design Wizard," shown in Figure 1. The wizard helps XPower users read in design data, input VCD files if available, set default parameters, and identify and then set specific input activity rates – all in a tab-delineated form that's easy and intuitive to use. This new wizard helps you get more accurate thermal estimates faster.

MXE-II Simulation Software

ModelSim™ Xilinx Edition II (MXE-II) simulation and debug software is now available to all ISE customers. Based on Model Technology's ModelSim 5.5e software, MXE-II delivers the additional simulation capacity and performance needed to verify the ever-increasing repertoire of faster and denser programmable devices. These additional capabilities are user-transparent, requiring no time-consuming learning curves. In fact, MXE-II simulation software requires less from the end user because of the integration with ISE's HDL Bencher™ graphical test bench generation environment. HDL Bencher waveforms are automatically translated into VHDL or Verilog, simulated with MXE-II, and the expected results can be back-annotated into the original waveforms. Discrepancies between expected and actual results are also highlighted, expediting dynamic verification.

Moreover, MXE-II includes optimized libraries that deliver faster simulation runtimes than competing technologies for post-route non-timing/timing simulation. These libraries are available for all the device families supported in ISE. These enhancements – coupled with a new, fully automated licensing process – significantly expand MXE-II's usefulness to a broader set of PC-based customers.

MXE-II performs 33% better than its predecessors. That, plus the doubling of MXE-II's capacity, makes it ideal for the verification of all types of Xilinx devices, including the CoolRunner-II, Spartan-IIE, and Virtex-II families.

Conclusion

ISE 4.2i continues to deliver the speed you need – from the simplest designs in the smallest device families to the most demanding high-speed, embedded system designs. Upgrade today to get the most leading edge capabilities available for your programmable logic products. You can find out more about ISE 4.2i at *www.xilinx.com/ise/42i*. For a recorded e-learning lecture on ISE 4.2i, go to *www.xilinx.com/support/training/north-america-home-page.htm.*

New Configuration Options for Virtex-II Pro

Configure Virtex-II Pro FPGAs – and load embedded processor software – using the System ACE pre-engineered configuration solutions.

by Frank L. Toth
Marketing Manager, Configuration Solutions Division
Xilinx Inc.
frank.toth@xilinx.com

Xilinx offers a variety of System ACE™ configuration options to meet a wide range of configuration speeds, densities (bitstream size), and costs, as shown in Table 1. These pre-engineered solutions simplify the configuration of our FPGAs and help you get your design to market as quickly as possible.

The System ACE CF solution, shown in Figure 1, leverages the tremendous density of commercial compact flash memory devices to give you the benefit of industry-wide increases in speed and density. When power is applied, the System ACE CF solution uses the JTAG port to configure the Virtex-II Pro™ fabric and the embedded processor, as shown in Figure 2. Virtex-II Pro FPGAs can also be easily accessed by the Xilinx ChipScope™ debugging tools using this same JTAG port.

The System ACE CF solution uses the industry standard FAT file management system, which enables you to manage both bitstream and software files like disk space. Multiple configuration bitstreams can be selected and managed under software and hardware control. Built In Self Test (BIST) files at the system level can be loaded and used to reconfigure the FPGAs to perform system-level I/O, network communication, functionality, and memory tests. Then the FPGA can be reconfigured for the mission mode.

Built-in Configuration Controller

In a system with a Virtex-II Pro device along with other Virtex™ and Spartan™-II FPGAs, the Virtex-II Pro device may be configured first. Then the embedded processor can be used as a configuration controller eliminating the need for a separate processor. The Virtex-II Pro FPGA

	System ACE CF (Compact Flash)	System ACE MPM (AMD Standard Flash)
Multiple Designs	no limit	up to eight
S/W Storage	Yes	No
Removable	Yes	No
IRL Hooks	Yes	Yes
Density	128Mbit - 8Gbit	16-64Mbit
Performance	~30Mbit/sec	152 Mbit/sec
#Components	2	1
Min. Board Space	~(50mm x 50mm)	~(35mm x 35mm)
Space/bit	32-256 Mbit/sq. in.	8-32 Mbit/sq. in.

Table 1 - System ACE configuration options

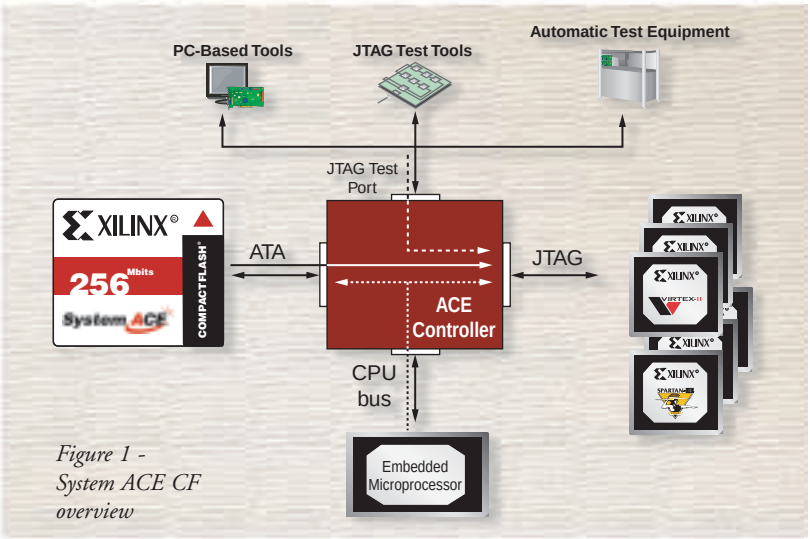

Figure 1 - System ACE CF overview

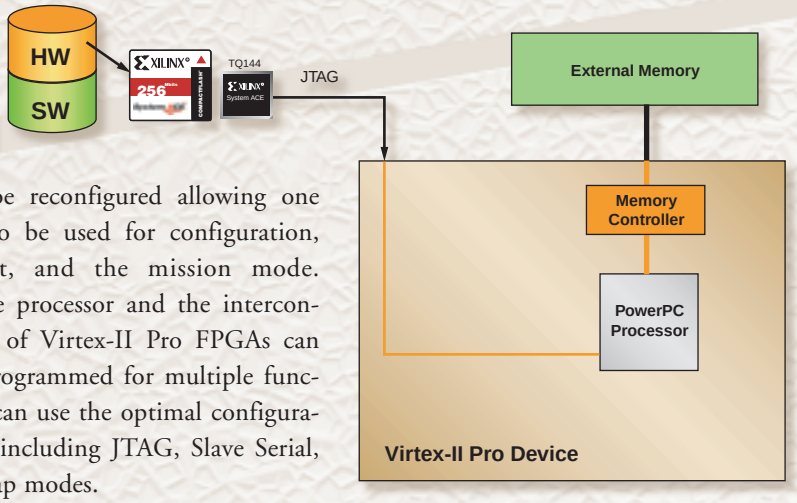

Figure 2 - Configuration via JTAG port

can then be reconfigured allowing one processor to be used for configuration, system test, and the mission mode. Because the processor and the interconnect fabric of Virtex-II Pro FPGAs can easily be programmed for multiple functions, you can use the optimal configuration mode including JTAG, Slave Serial, or SelectMap modes.

As an alternative configuration sequence, the Virtex-II Pro FPGA can be "booted," as shown in Figure 3, using a smaller boot file provided by the System ACE solution on power up. The embedded processor can then take control of the configuration process and in turn configure itself and the rest of the system by sending the appropriate interface and control signals to the System ACE solution.

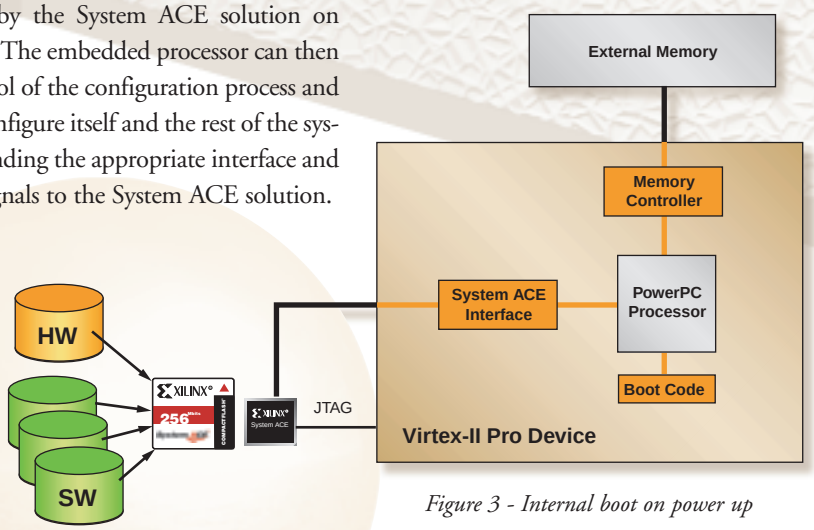

Figure 3 - Internal boot on power up

Other Configuration Options

Xilinx also offers several other pre-engineered configuration solutions including serial PROMs, the System ACE MPM (Multi-Package Module, shown in Figure 4) and the System ACE SC (Soft Controller). The density of the System ACE MPM solution ranges from 16 to 64 megabits (or more by using the System ACE bitstream compression software) and features a Slave Serial or Select Map configuration port output. The native flash/microprocessor interface allows the System ACE MPM solution to be connected to the microprocessor bus.

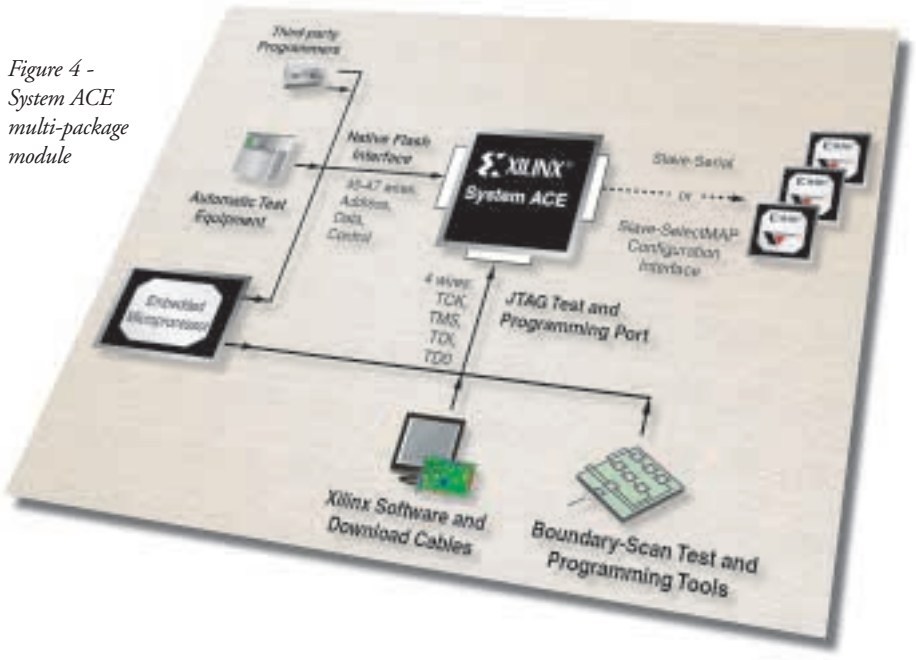

Figure 4 - System ACE multi-package module

If you want to use a different form factor and mount the components directly on the board, the System ACE Soft Controller IP is available. This pre-engineered solution is identical in functionality to the System ACE MPM solution and you can download the IP free of charge.

Taking Full Advantage of System Re-Configurability

With the advent of the embedded processor and the other system-level features of the Virtex-II Pro family, the use of multiple configurations becomes more useful, both to increase system flexibility and extend product life. Multiple configuration files can also be used as a vehicle for loading BIST and for testing system functionality, integrity, and performance.

Conclusion

With the Xilinx System ACE solutions, you can easily control both the configuration of your FPGAs and automatically load your embedded processor software. There is no easier way to configure your Virtex-II Pro designs.

For more information on System ACE products, go to: *www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=system_ace*

Want to pump up your FPGA processing performance?

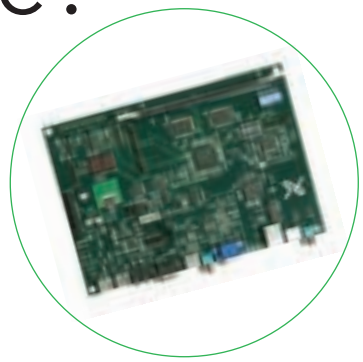

We can show you how we transformed a soft 32-bit microprocessor core running at 50 MHz in a Virtex-II FPGA into a cruncher that can do operations that would challenge GHz microprocessors.

If you want to do real-time video manipulation and storage while processing audio and handling the latest networking protocols — add "Xyronium™" technology to your multi-tasking embedded control environment and realize incredible performance gains.

Need to do some heavy duty processing? Order one of our development boards today to realize your full potential.

Find out how Xyron unleashes the untapped power of today's processors.

Visit us at Embedded Systems Conference – San Francisco Booth 639s or at "Xilinx Programmable World 2002", or see us online at **www.xyronsemi.com**.

xyron™ semiconductor

www.xyronsemi.com

See us at the Xilinx Programmable World 2002

Platform FPGAs Take on ASIC SOCs

Here are seven good reasons why Platform FPGAs provide a superior design environment and faster time to market than ASIC SOCs.

by Milan Saini
Technical Marketing Manager,
Alliance Software Marketing
Xilinx, Inc.
milan@xilinx.com

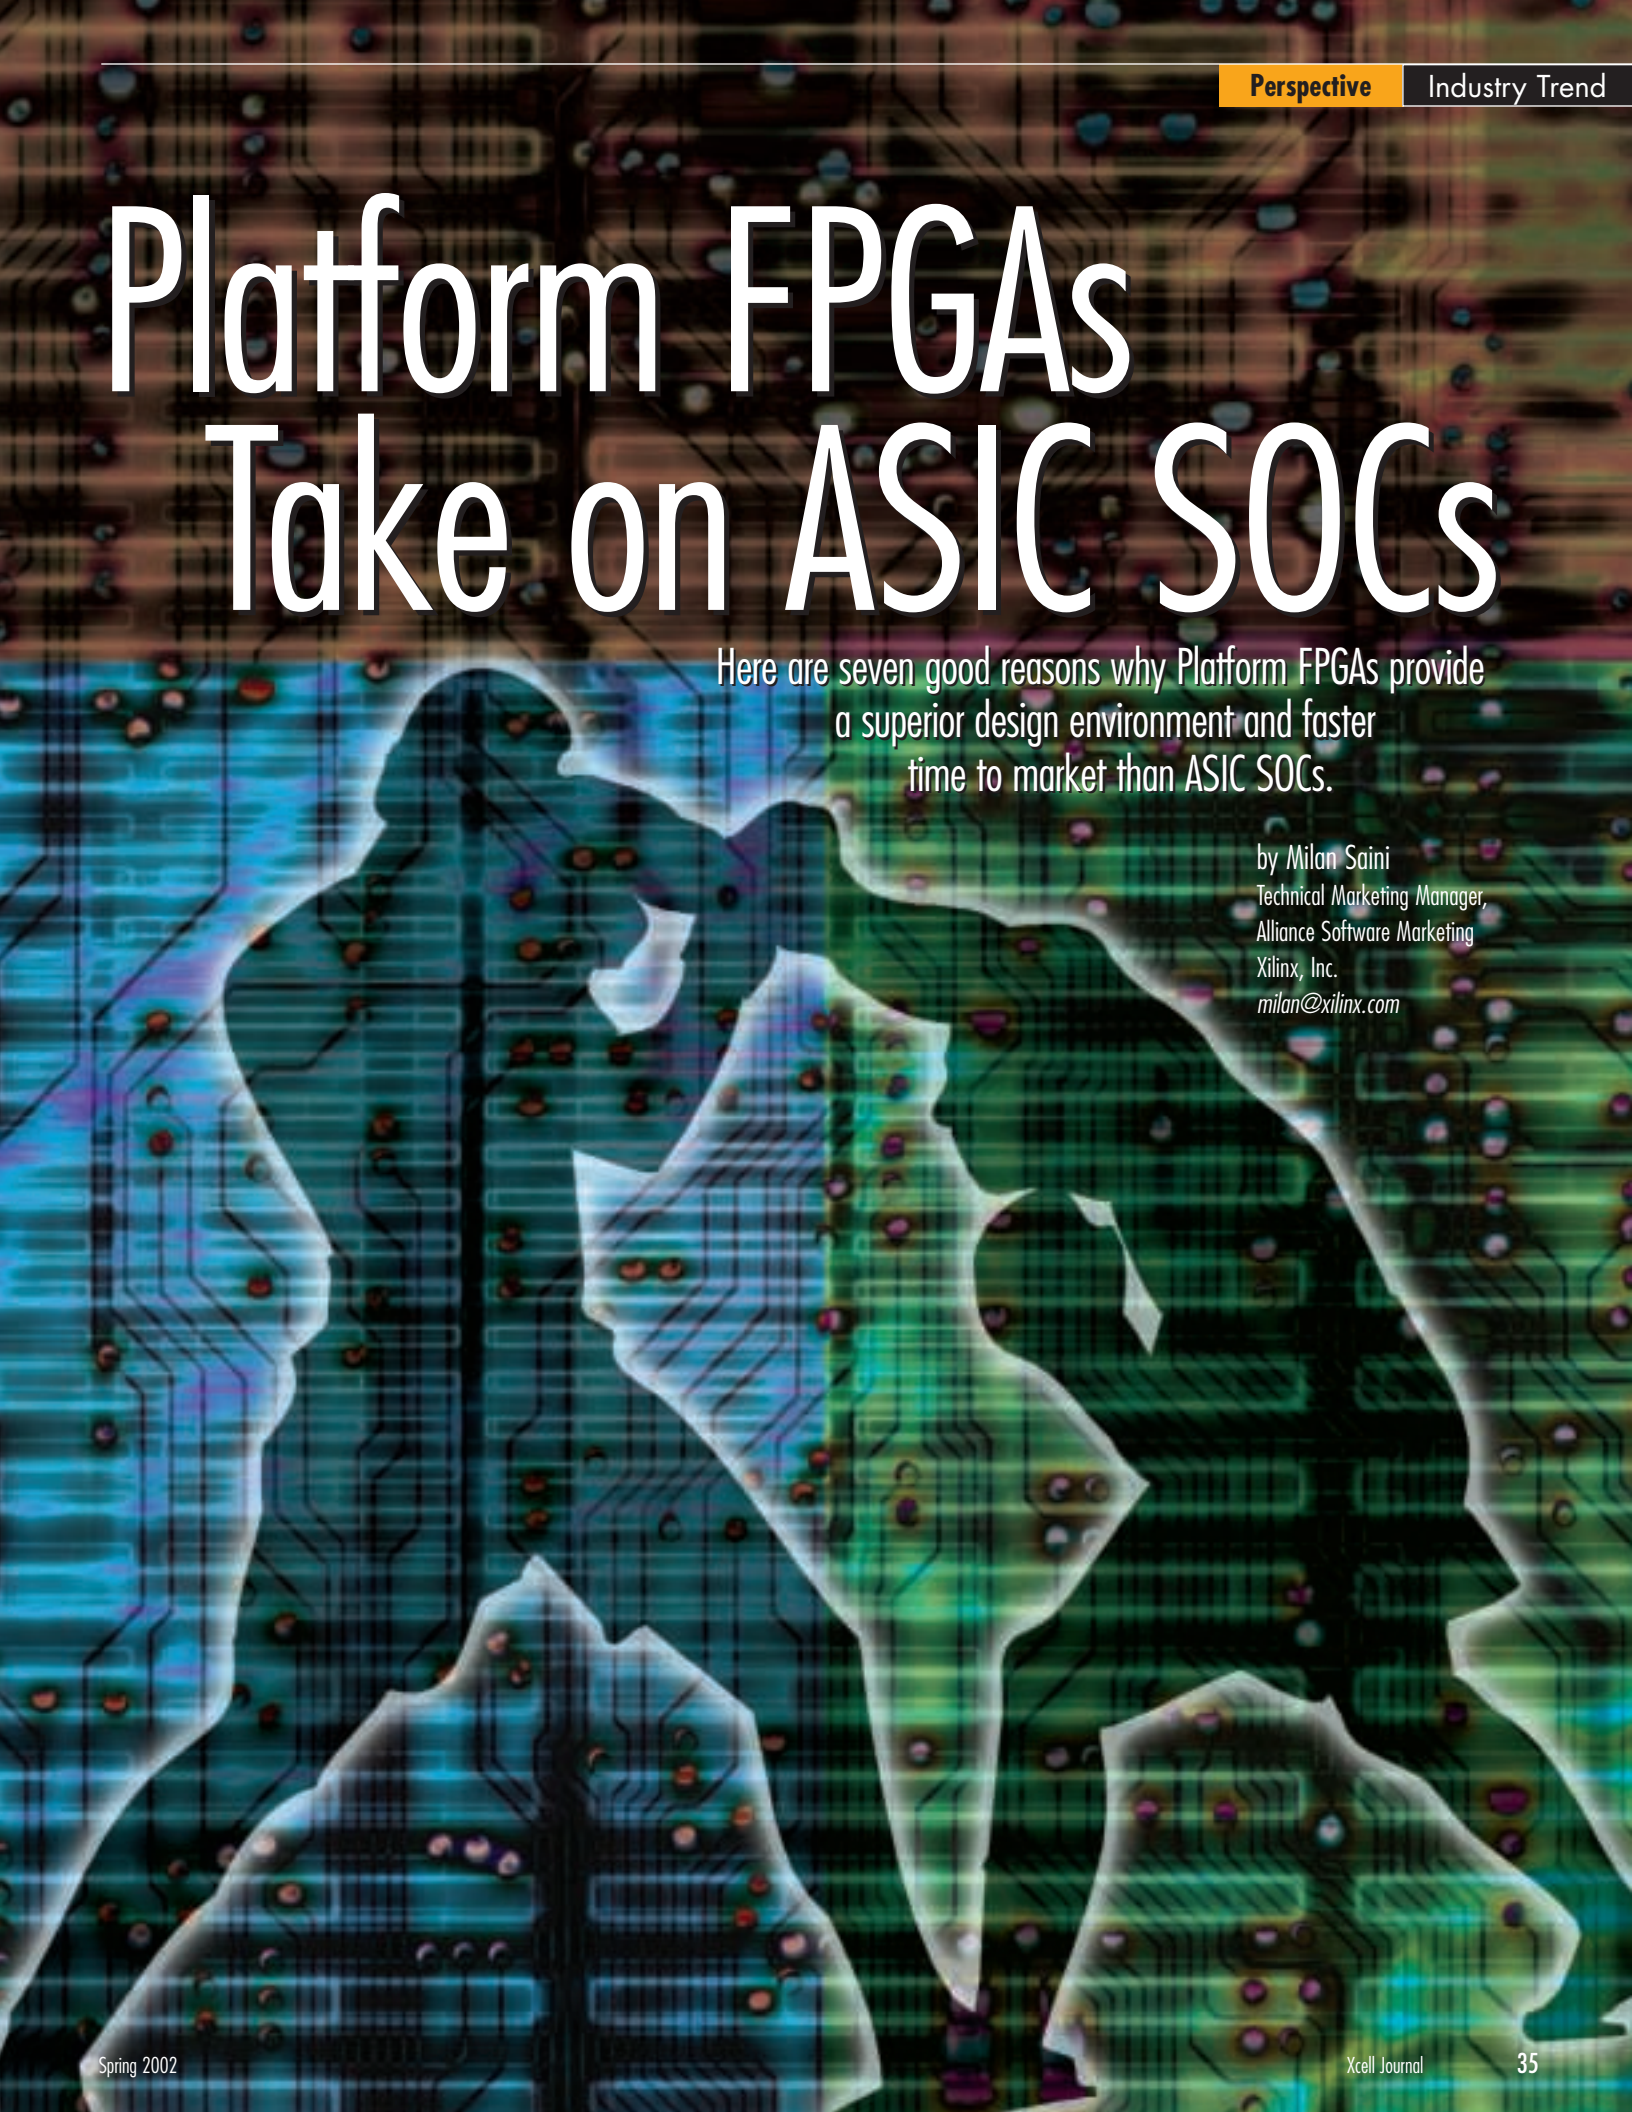

The system-on-a-chip (SOC) market has experienced steady and consistent growth. Dataquest estimates roughly half of all ASIC design starts are SOC based. That percentage is projected to reach 80% by 2005. There are several reasons for this clear shift in design methodology. Some of the obvious advantages include greater component integration, increased speed (Logic <-> Processor), lower packaging and test costs, and increased overall system reliability. All of these combined can potentially make significant contributions to achieving the often elusive but always important goal of accelerated time to market.

Programmable logic device vendors have entered the SOC solution space with the introduction of a new class of devices known as Platform FPGAs – with the most recent addition being the Virtex-II PRO™. These devices offer the same level of integration as ASIC SOCs, but in contrast, Platform FPGAs facilitate the development of a wide range of applications on the same chip. This article focuses on seven of the key advantages of the Platform FPGA approach.

#1 – Pre-Engineered Platform

Platform FPGAs integrate several fixed and predetermined blocks of hard IP (intellectual property) components (system elements) within the programmable fabric. Notable among these are high-performance RISC CPUs, multi-gigabit and high speed I/Os, block RAM, system clock management, and dedicated DSP processing hardware. This powerful assembly and harmonious blend of components creates a cohesive system design environment. This environment offers unprecedented flexibility and performance, thus enabling the deployment of a wide range of applications.

The critical technological breakthrough is in the ability to tightly interface the various elements into the programmable fabric. Without this tight integration, much of the speed benefits could not be realized. The fact that the choice of system elements is already made greatly simplifies the design and development process. A fixed architecture is particularly beneficial for software tool and IP providers in allowing them to deliver better value, customization, and architecture-optimized solutions.

The assembly of the various hybrid IP blocks in an ASIC adds substantial complexity and hardship to the users' design and development environment, because of a variety of issues relating to tool and IP interoperability, physical layout, timing, and system verification. For Platform FPGAs, on the other hand, it is much easier to tailor and optimize components – such as silicon, software, support, and IP – because FPGAs represent a fixed and pre-engineered target.

Summary: A fixed, pre-engineered but programmable FPGA solution offers a more productive and efficient development environment from both the software and silicon perspective.

#2 – Process Technology

PLD vendors have been able to extract great value and benefits from Moore's Law and shrinking device geometries. While the majority of ASIC design starts are at or higher than 180 nm, FPGAs have raced ahead to bring the cost and performance advantages of 130 nm to its customers. This ability to rapidly migrate to the leading edge of technology is essential to delivering the performance, capacity, and integration that is necessary to challenge and displace the current established platforms of system design.

For instance, FPGAs now make it entirely feasible to build systems of up to 2M (ASIC) gates with CPU(s) running at 400 MHz, serial I/O channels at 3.125 Gbps CLB fabric-switching at 300 MHz, and the entire system clocking over 150 MHz. This surpasses the projected sweet spot of ASIC SOC designs.

Looking at it in pure economic terms, the costs for a typical mask set for a 130 nm ASIC run into the $700K+ range. That raises the bar for entry into the ASIC space, making the Platform FPGA an even more attractive option for a growing percentage of all SOC designs.

Summary: FPGA silicon is best-in-class in process and engineering, thereby delivering best-in-class system performance.

#3 – Software Tools and Methodology

It is well known that EDA designer productivity for ASICs is lagging behind recent silicon advances. The ability to create a productive design and debug environment is absolutely essential to the success of any silicon platform. Therefore, software plays a critical role in not only making the platform easy to use and work with, but also in extracting the most performance and device utilization out of the silicon.

The goal is to insulate the user from having to learn extraneous details about the platform, yet providing empowerment and control when and where it is needed. To this end, FPGA vendors have created customized and user-friendly processor system generator tools that aid in instantiation, initialization, and configuration of all the various system component blocks. In addition, these software tools automate otherwise manual and error-prone tasks, such as the interconnections among the processor, its peripherals, and buses.

Design entry is engineered to tightly couple the HW/SW domains. Such engineer-

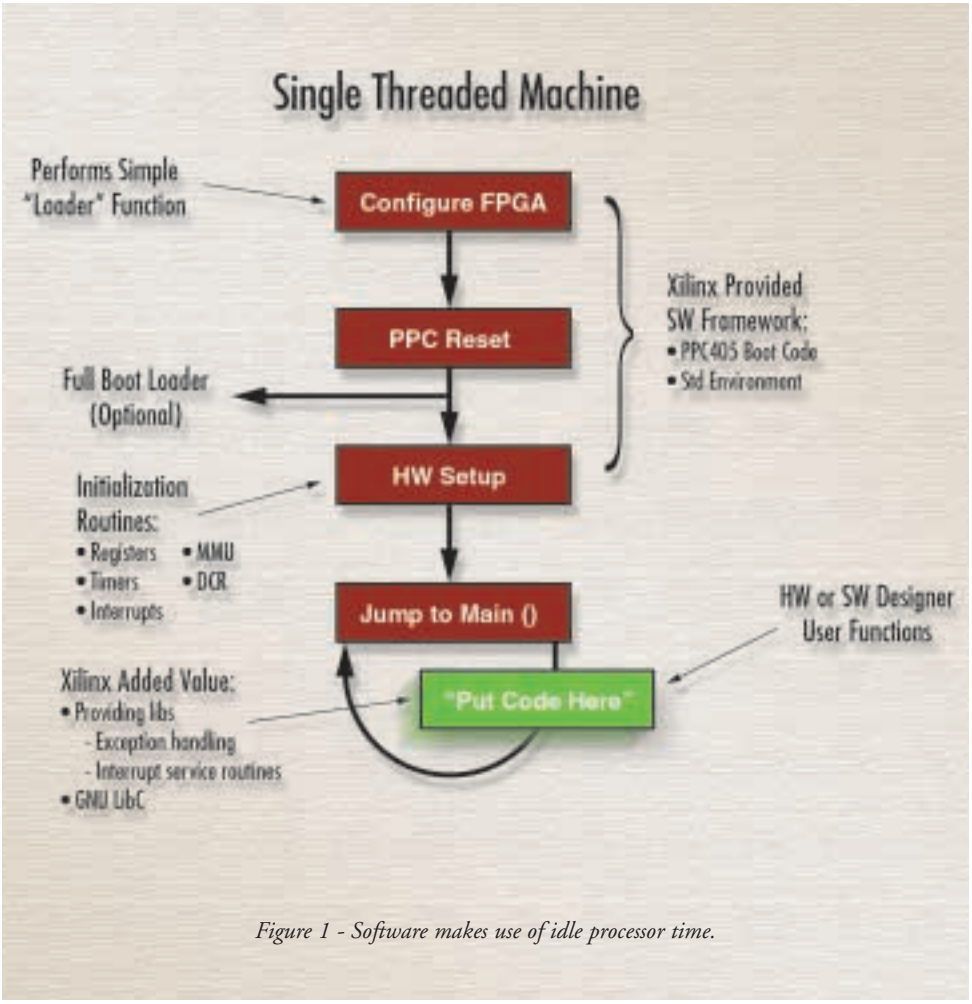

Figure 1 - Software makes use of idle processor time.

ing leads to not only simplified and easy-to-use design flows, but it ensures reliability and robustness from the very start by performing design rule and data consistency checks across the two domains.

Two examples of the advantages of cross-domain HW/SW co-design include:

- Automatic generation of device drivers and header files for SW engineers once a particular block is instantiated by the HW designer
- Tools to automatically populate SW binary code into appropriate FPGA memory bitstreams.

Summary: Software sells silicon.

#4 – Advanced Debug

The importance of finding problems early cannot be overstated. Yet, up-front ASIC verification is extremely designer- and computer-resource intensive. Compared to systems on a board, SOC limits visibility into the internal nodes of the system, making the task of verification and debug more challenging than ever.

The critical part includes the verification of complex interactions between the application software and the custom-designed peripheral hardware. Traditional HDL-centric verification and debug techniques can no longer deal with the rising complexity of system designs.

Consider a typical application, such as MPEG A/V decoding, where a large number of simulation cycles are required to complete a small sequence of frames. Co-verification tools in this case would either take an impractically long time to complete or validate only a mere fraction of the software code, falling far short of what it takes to find problems in the HW/SW interface.

FPGAs overcome this problem in large part by being able to provide access to real or near real targets at a very early stage in the design cycle. Among other things, this means that SW engineers using FPGAs can quickly and easily sort out logic and design flaws by targeting real silicon. The engineers do not have to rely on inefficient ASIC-centric techniques like co-verification or writing stub code. Application software can be debugged at system speeds with full hardware and software register access and control.

Additionally, the FPGA fabric allows for construction of highly customized and value-added cores to enable powerful real-time, on-chip debug capability. Some examples of such instrumentation include:

- Logic and bus analyzer functions
- Bus protocol compliance monitors
- Memory buffers for debug and trace port data
- Cross-domain triggers and breakpoints
- Hardware run control of the CPU
- HW/SW time synchronization logic.

Furthermore, a single cable is able to perform multiple functions, like debugging hardware, debugging software, as well as programming the FPGAs. This greatly simplifies the lab setup making it much easier to exploit the debug advantages.

Summary: Platform FPGAs offer a clearer, more cohesive, and overall more effective debug strategy. Specifically, Platform FPGAs offer up-front silicon access along with unprecedented visibility and control of the processor and its peripherals residing in the programmable fabric.

#5 – Top Tier Partnerships and Vendor Tools Support

In extending the concept of traditional programmable logic to Platform FPGAs, certain critical technologies have had to be developed or acquired. One of most exciting aspects brought forth by FPGA vendors has been to successfully forge

strategic partnerships with vendors holding various system technology components. Driven largely by the successes of the FPGA business models, leading vendors have been eager to partner in fulfilling the vision of building powerful programmable systems platforms.

Areas of cooperation include partnerships in the form of cutting-edge process technology, powerful mainstream processing elements – such as RISC CPU, high-speed I/O – and other components deemed of significance to a system design platform. IBM, Conexant Systems, Inc., Wind River Systems, and other high profile vendors are currently engaged in such strategic partnerships. What this means to you is there is no need for these partners to negotiate licensing, royalty, and integration issues with individual vendors, thereby greatly reducing your engineering, management, and accounting overhead.

The appeal and draw of FPGAs has caught the attention of independent SW tool vendors. Increasing numbers of vendors are able to sustain business models selling to FPGA customers. Several new vendors are setting up shop to write custom tools to help enhance and exploit the unique capabilities of Platform FPGAs.

Summary: The FPGA business model has attracted top-tier silicon and IP vendors to forge strategic partnerships to create powerful system design platforms. Software vendors are able to financially justify investment in research and development leading to a continuous stream of an increasing number of innovative solutions.

#6 – Application Space: Co-Design Flexibility

Programmable HW combined with processors on a single chip softens the HW/SW design boundary. By using hardware-assisted architectural exploration, designers can optimally search for the right HW and SW partitioning, which leads to the increased probability of being able to meet performance and area targets. Sequential computing, exception handling, and control functions, for instance, programmed in HW could be implemented in SW running on the processor to save silicon.

On the other hand, SW structures and algorithms – which can be broken into parallel, non-blocking processes – can achieve significant speed and data throughput improvements by implementing them in HW. In fact, software engineers represent a new and emerging market for Platform FPGAs. Aided by design tools, it is now easier than ever for SW engineers to explore concepts of software acceleration via HW. Both hardware and firmware are reprogrammable and field upgradeable, which enables the development and deployment of several product generations from one base. This translates into a much broader applicability than ASICs and ASSPs.

In addition to being suitable for building complex embedded applications, HW engineers can utilize an otherwise idle processor to run relevant portions of their design algorithms or control logic. As Figure 1 shows, the CPU can serve as a simple microcontroller running a single-threaded application. PLD vendors add value by providing software device drivers and library functions for rapid implementation without requiring the HW engineer to have detailed knowledge of SW practices.

Summary: The Platform FPGAs provide the most flexible co-design platform by enabling dynamic HW-SW design partitioning.

#7 – Risk Management

When compared to Platform FPGAs, ASIC SOCs present a huge design risk. With more variables and issues to worry about, and with limited debug capabilities when mistakes are found, ASIC designers are forced to either resolve problems suboptimally in software, or in the worst case scenario, they are forced to respin the silicon at great cost and loss of time to market.

Reprogrammability comes up big here. Programmable platforms allow early access to silicon. Engineers can validate performance and functionality in real silicon, leading to greater confidence in the reliability of the completed system. The programmability of the platform allows itself to be debugged and upgraded even after the system has been deployed. This helps promote and protect the time-to-market advantage by alleviating a large part of the risk.

Summary: Programmable Platform FPGAs provide better control over the life cycle management of products by minimizing the cost and time penalty for design errors and specification changes.

Conclusion

In an era when SOCs continue to dominate mainstream design, Platform FPGAs are emerging to take a share of the spotlight from ASICs. While ASIC SOCs have and will continue to be strong in certain segments – like low power, small form factor, and high-volume, cost-sensitive consumer electronic gadgets – an increasing range of other infrastructure applications in areas such as networking, telecommunications, industrial electronics, and data storage have compelling reasons to move to a programmable platform.

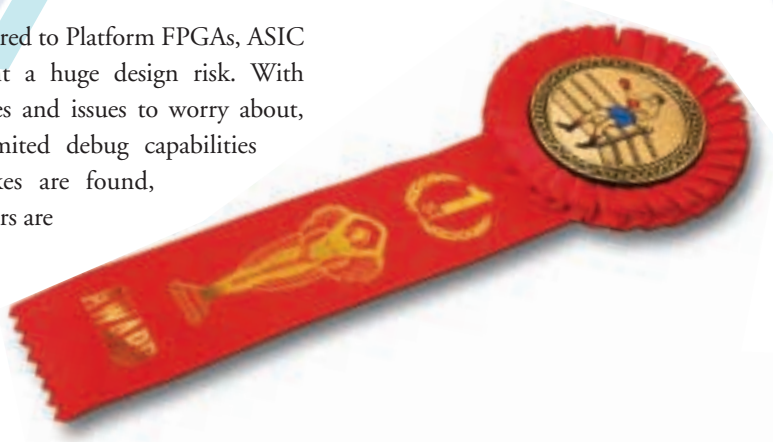

FLEXTRONICS®
Design

Flextronics Design, a business unit of Flextronics, is a leading provider of engineering design services delivering complete system-level solutions.

With 25 locations globally and more than1200 design engineers, Flextronics Design offers expertise in:

-FPGA and SOC ASIC Design
-Embedded System Design (Hardware and Software)
-PCB Layout
-RF Design
-Industrial and Mechanical Design
-Tooling
-Test Development
-Compliance & Regulatory Test Engineering

Flextronics Design provides the technical expertise, engineering depth and design capability to handle every phase of the product development process, from conception to production launch.

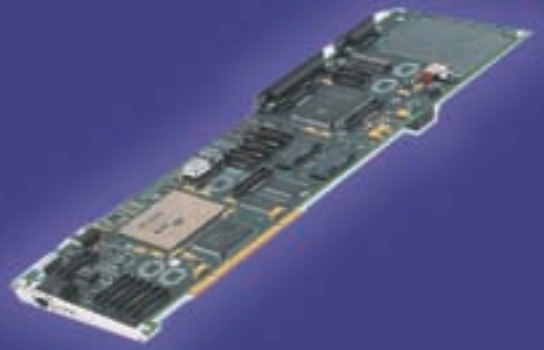

Visit us on the Web at: www.flextronics.com

See us at the Xilinx Programmable World 2002

Your Reconfiguration Is in the E-Mail

With Xilinx Internet Reconfigurable Logic technology and Virtex Platform FPGAs, you can perform fast and easy remote field upgrades via e-mail using microcontrollers.

by Marc Defossez
Senior Staff Applications Engineer
Xilinx, Inc
marc.defossez@xilinx.com

Since the beginning of the FPGA technology, Xilinx has pushed the boundaries of reconfiguration. In earlier FPGA families, it was only possible to reconfigure the whole FPGA. With the introduction of the Virtex™ FPGA families, it became possible to partially configure or partially reconfigure an FPGA. It is also now possible to reconfigure a remote FPGA via the Internet using Xilinx Internet Reconfigurable Logic (IRL™) technology. However, only a few companies and a few of all FPGA designs make use of IRL technology, because of the perception it is expensive, complicated, and mostly a proprietary solution.

What if we could securely reconfigure FPGAs in the field simply by sending an e-mail message? In this article, we will show you just how easy and cost-effective that can be.

Protocol Stack

Xilinx IRL reconfiguration technology uses the same transmission protocols as everyday Internet e-mail:

- **TCP/IP** – Transmission Control Protocol over Internet Protocol transports the e-mail over the Internet to its destination.
- **SMTP** – Simple Mail Transfer Protocol is used to deliver the messages.
- **POP3** – Post Office Protocol 3 retrieves the messages.

Figure 1 shows a complete Internet protocol stack. Each layer of the protocol stack is an abstraction level hiding details from other layers on top or below. For example, the network access layer does not need to know what kind of data it is carrying. Whether the data is video, voice, or other, it is unimportant to the network access layer. The only thing this layer needs to do is deliver the data in good quality to the upper layers.

- **Application Layer** – This is the user interface layer. The POP3 and SMTP protocols necessary for IRL technology to work reside in this layer.
- **Transport Layer** – This layer implements reliable, full-duplex communication over the Internet. TCP works in this layer.
- **Internet Layer** – Addressing and routing of data happens in the Internet Layer, where IP is implemented.
- **Network Access Layer** – Also called the "link layer," this is where the hardware interface is managed.
- **Physical Layer** – This is the actual medium for communication across great distances. Such transmission media include co-axial cable, phone lines, fiber optics, and wireless broadcasting.

Implementation

If you want to implement the Internet stack into an FPGA as hardware, it will:

- Take a lot of time (including VHDL or Verilog design and simulation).
- Require a robust FPGA.
- Consume a lot of money.

On the other hand, microcontrollers are good for protocol handling and can mitigate the time and cost of building an IRL solution for the remote reconfiguration of FPGAs.

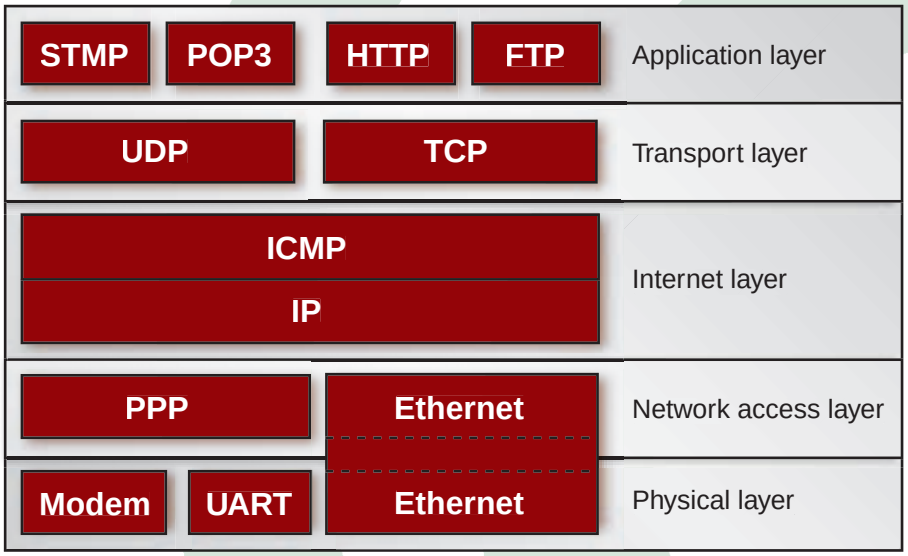

Figure 1 - Internet protocol stack

Two microcontroller solutions are possible:

- Use external microcontrollers.
- Put a microcontroller inside a Virtex Platform FPGA.

There are two ways to embed a microcontroller in a Virtex device.

- You can use the software MicroBlaze™ microcontroller in Virtex, Virtex-E, Virtex-II, or the new Virtex-II PRO™ Platform FPGAs.
- You can buy a Virtex-II PRO Platform FPGA with a hard-wired IBM® PowerPC™ 405 microcontroller already onboard.

External Microcontrollers

Several microcontroller manufacturers have Internet capable components. For instance, two Ubicom microcontrollers – SX52BD and IP2022 – can be used for IRL applications. Such external microcontrollers require "virtual peripherals" from Ubicom and some small modifications and additions to control downloading to an FPGA.

These small controllers and peripherals make the implementation of the TCP/IP, SMTP, and POP3 components of IRL technology easy and fast. Due to the small amount of internal memory of the microcontrollers, however, the Internet protocol stack is tuned to only perform the necessary functions.

Basic Setup

The simplest setup consists of an FPGA and a small controller, as shown in Figure 2. Here's what happens when the system is powered up:

- The FPGA is empty.
- The microcontroller waits for a certain time until all components of the IRL design have reached a stable state.
- Then the controller connects to the network by sending AT commands to an external modem through an RS-232 device, or by sending AT commands to an onboard chip modem, or by other physical implementation.
- Once the link has been set up with the e-mail server, the microcontroller asks if there is e-mail available. When there is, the microcontroller checks the e-mail header.
- If the header is not of a particular type, the controller will delete the mail message on the server.
- When the e-mail has the correct header type, the microcontroller downloads it. The /PROGRAM pin is toggled, the contents are serialized onto an output pin, and a bit clock is generated.
- When the DONE pin goes High, the microcontroller deletes the e-mail on the server.
- The microcontroller breaks the connection.

If the DONE pin does not go High after a period of time, the download operation is repeated until the DONE pin goes High.

Although this is the simplest approach, it has its shortcomings:

- An Internet connection is obligatory.
- The server must always have mail ready for the application, or else the application cannot start.
- When configuration fails, no fail-safe recovery mechanism exists.
- The design can go into an endless loop trying to download its configuration.

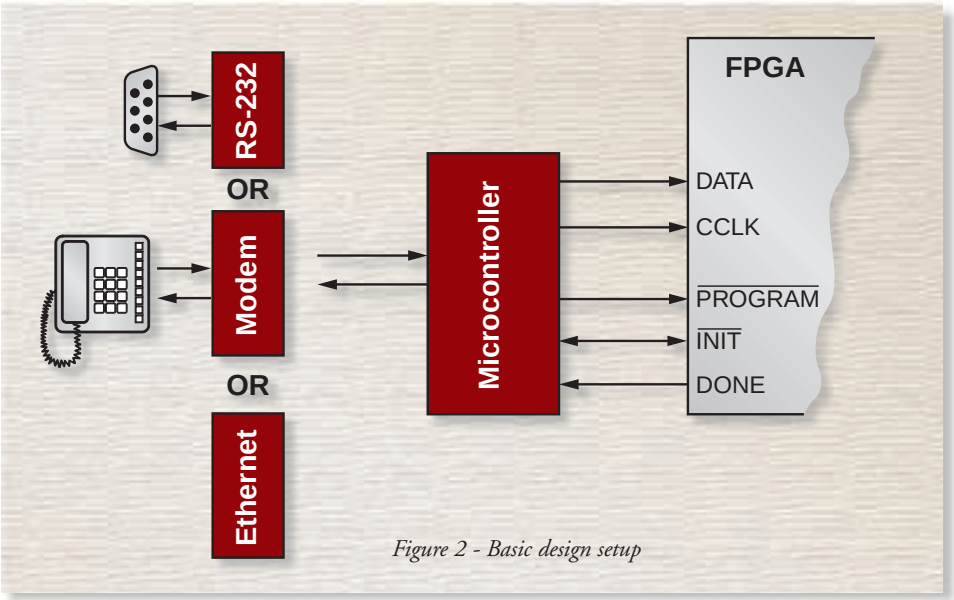

Figure 2 - Basic design setup

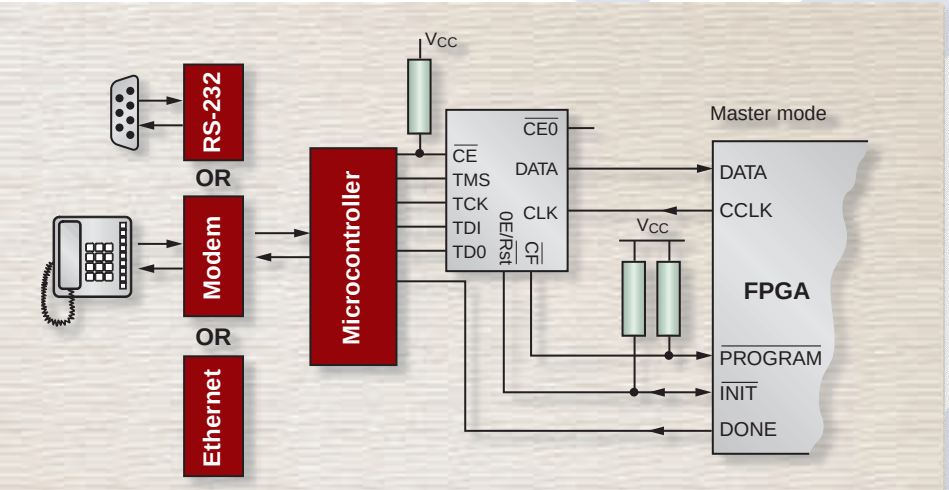

Figure 3 - Safe download

Fail-Safe Setup

You can make the download more reliable by storing the downloaded bitstream into semipermanent memory (flash RAM). The FPGA can then be reconfigured from the flash memory. See Figure 3.

An even more secure solution is to work with two memories. A basic configuration can be loaded into the FPGA when it is shipped from the manufacturer. During operation in the field, the microcontroller can connect to the Internet and download a new configuration into the second memory. The new configuration bitstream would be downloaded into the FPGA at next boot.

When the download works, the new configuration will be used. If the new programming bitstream fails, the microcontroller will boot again from original memory.

Internal Microcontrollers

Internal microcontrollers for FPGAs come in the form of MicroBlaze software and the hardware-embedded PowerPC 405 processor in Virtex-II PRO devices. (As a point of interest, the Virtex-II PRO Platform FPGA can be configured with both the MicroBlaze software and the PowerPC 405 processor, but that is beyond the scope of this article.)

While there are significant advantages of having microcontrollers onboard Virtex Platform FPGAs, there are also factors, in this design case, that require special consideration:

- In both cases of soft and hard microcontrollers, the Internet protocol stack must be programmed (ported) onto the FPGA.
- The FPGA must get a basic (possible partial) bitstream to download the MicroBlaze controller, its memory, and peripherals.
- The PowerPC Virtex-II PRO Platform FPGA must have memory and peripherals downloaded before it's able to boot.
- The small control algorithm that was done in the previous description within the external microcontrollers must now be implemented in a small CoolRunner™ CPLD, as shown in Figure 4.

Figure 4 - Internal processor

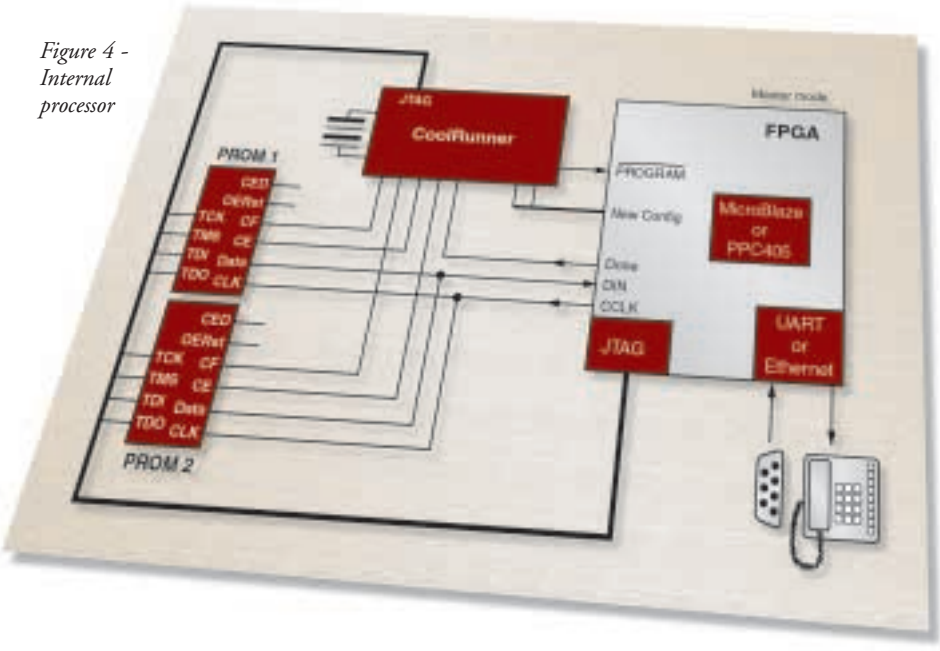

In either case, the first bitstream must contain the basic application of the FPGA. This way the system can operate as a standalone unit without problems. However, during operation, the FPGA microcontroller can contact the Internet and download new bitstreams as they become available.

The "basic setup" for external microcontrollers described above cannot be applied to internal microcontrollers, because the downloaded bitstream must be stored by the FPGA in some semipermanent (flash) memory. The final design of embedded microcontrollers can have different levels of fail-safe operation, depending on the system requirements.

Conclusion

It is possible to perform IRL operations in a safe and inexpensive way using external or internal microcontrollers – and some extra hardware. Xilinx IRL technology can now be used in almost any design that requires field upgradability. Reconfiguration by e-mail eliminates the need for complex and expensive design solutions. Furthermore, with a little ingenuity, the IRL solution can be extended to designs that were not previously considered for field upgradability. Application notes with detailed solutions will be posted in the coming months on the Xilinx website: *www.xilinx.com/apps/appsweb.htm.* Look for "Virtex FPGA" or "FPGA Configuration."

the only logical solution

NALLATECH

SPEED
POWER
FLEXIBILITY

Securing a competitive edge in a programmable world requires shorter development times, significant reduction of risk and cost, and seamless integration of products with target systems. For over 10 years Nallatech design engineers have delivered all these advantages of generic reconfigurable computers to global OEMs in 21 countries.

www.nallatech.com

See us at the Xilinx Programmable World 2002

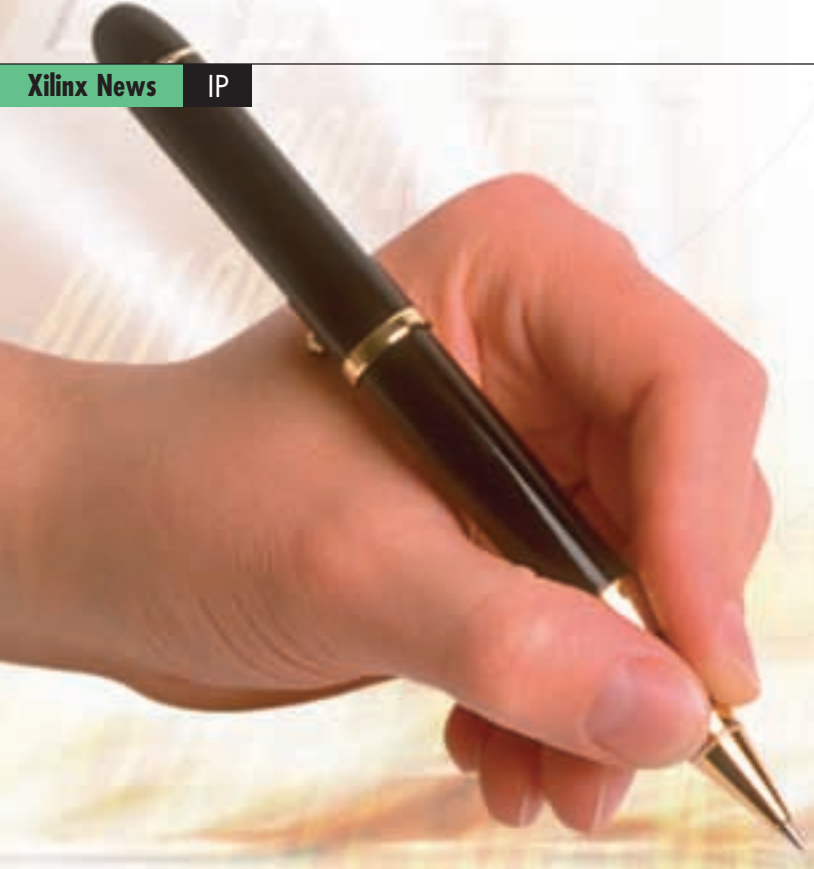

SignOnce and Break the IP License Barrier

Xilinx has sponsored the Common License Consortium to streamline the IP licensing process and improve your time to market.

by Mark Bowlby
AllianceCORE Program Manager
Xilinx, Inc.
mark.bowlby@xilinx.com

As a designer, you are constantly asked to deliver new and enhanced systems. Your marketing department asks you to add new features, but do you have the time or prior experience to design everything in the new specification? If not, you can always look at including pre-developed system functions in the form of intellectual property (IP) cores from the FGPA vendor or one of a growing number of third-party providers. But can you wait until the legal issues around IP licensing are resolved before you even begin your design? If only there was a way to clear the legal hurdles so you could spend your time on what matters most – getting your design completed on time.

In September of 2001, Xilinx launched the Common License Consortium to specifically address this issue. It is an industry-first initiative to simplify the licensing process for FPGA-based IP cores. At that time, Xilinx and more than 21 third-party IP providers all agreed to offer a common set of IP licensing terms – called the SignOnce IP License – to get access to use their cores. Let's look at how you can benefit from this new and innovative program.

Strength in Numbers

When selecting IP cores for your FPGA designs, you have access to Xilinx LogiCORE™ products as well as products from our third-party network of AllianceCORE™ providers. This hybrid approach ensures you a broad portfolio of IP and expertise to choose from. You benefit from competition, and it potentially puts experts right in your own backyard.

But everything comes at a price. A multi-vendor solution introduces the challenge of dealing with multiple providers, each with a different set of IP license terms. Research has shown that license negotiations between a customer and a single supplier can exceed six months. Clearly, situations that require you to deal with more than one vendor would be intolerable.

The SignOnce IP License takes care of this. With it, you can sign up to a single set of license terms that, today, gives you access to 500 IP cores from well over 30 providers. Consortium membership continues to grow and includes companies from North America, Europe, Japan, and Southeast Asia, so there is likely a solution provider somewhere near you. There is no cost for membership in the consortium, and only IP providers may join. They do not need to participate in the AllianceCORE program to become a member, and there is no cost to customers to gain the benefits of the program.

What Is Covered?

The SignOnce IP License leverages the fact that licensing IP cores for use in a specific FPGA is less complex than licensing for use in an ASIC. ASIC use usually requires the transfer of source code. This requires extensive legal wording that can swell to 30 pages to protect the vendor against such issues as improper use and piracy.

The SignOnce IP License reduces this page count to four by focusing on the transfer of an FPGA netlist. Netlists are specific to particular FPGAs, and they are difficult to reverse engineer or port to a different technology. This dramatically simplifies the license process. After a SignOnce IP License is in place, you will then be able to purchase individual cores at a price that you and the vendor(s) agree on.

Usage Options

The SignOnce IP License provides project-based and site-based usage options.

A project-based license allows you to use a core within a single design. Ultimately, the core is incorporated into a larger design, converted to a bitstream, and programmed into the FPGA. Based on this, the definition of a "project" includes the following scenarios:

- A single bitstream, which can include multiple instances of the core. You can then use that bitstream on one or more printed circuit boards. In this case, the project is defined as the entire chip that includes the core, which can then be used (without modification) in other designs. This is similar to the way you might use an ASIC or ASSP device.
- A single printed circuit board, using one or more bitstreams, each containing one or more instances of the core. This allows you to leverage the reconfigurable nature of FPGAs in your design.

Your design group for the "project" can span multiple sites. Usage of the IP core that goes outside of the above definitions would require you to negotiate additional fees with the IP vendor.

On the other hand, a site-based license allows your company to use the core in unlimited designs developed at the specified site. You would have to pay additional fees for usage at other sites, or you could negotiate a list of sites to be covered under the up-front license fees.

In general, site-based licenses cost more than project-based licenses from IP vendors that support both. Most consortium participants support both license types, but some only offer one. You should inquire about this when working with the vendors.

Simplifying the Process

The SignOnce IP License consists of three parts: the legal terms contained in the body plus two exhibits.

The legal terms form the bulk of the license and govern issues such as usage (project versus site), intellectual property rights, indemnity, warranty, export restrictions, governmental use, and so on. All legal terms (except for usage) are identical for both project and site versions of the license.

Figure 1 - IP Center Smart Search with SignOnce IP license option

Exhibit A is designed to list the specific cores that you will license. Once you have a SignOnce IP License in place with a consortium member, you can license additional cores from them by negotiating a separate Exhibit A for each core and paying the appropriate license fees.

Exhibit B is designed for listing the consortium members that you wish to do business with. Each member signs and attaches a separate Exhibit B so you can add IP providers at will. You can sign all members on at once or add ones as needed by simply having each sign a separate Exhibit B that references the original license terms.

This structure provides considerable flexibility. Even if you sign a license with one vendor for one core, you have the ability to later sign on additional consortium members to purchase more cores. Because members have agreed to use the same licensing terms, additions take little further assistance from your legal department. When dealing with any consortium member, make sure you let them know that you are interested in licensing IP using the SignOnce IP License.

In Search of SignOnce IP

All SignOnce IP cores available from Xilinx and our AllianceCORE partners can be found in the Xilinx IP Center (*www.xilinx.com/ipcenter/*). As shown in Figure 1, the Smart Search engine allows you to restrict your searches to only show cores that are available under the SignOnce IP License. Even if you don't select this option, the search results will indicate which are SignOnce cores. You will need to contact non-AllianceCORE members of the consortium directly to find out what cores they offer.

Conclusion

Your job as a designer is already difficult enough. IP cores are available to simplify and accelerate your development process, allowing you to focus on the portions of the design where your expertise adds value. If you adopt the SignOnce IP License and deal with members of the Common License Consortium for your IP, you will remove a major time-to-market bottleneck for you and your company.

For more information on the SignOnce program – including the growing list of consortium members, copies of the licenses, and searchable lists of Xilinx and AllianceCORE vendor IP – follow the Web links shown in Table 1. The solution is available. Let it work for you.

Information	Web URL
Program information, consortium members, and license forms	*www.xilinx.com/ipcenter/signonce.htm*
Xilinx IP Center Smart Search Engine	*www.xilinx.com/search/ipsearch.htm*

Table 1 - Web links for SignOnce program information and products

Bring on the Music—But Take out the Noise

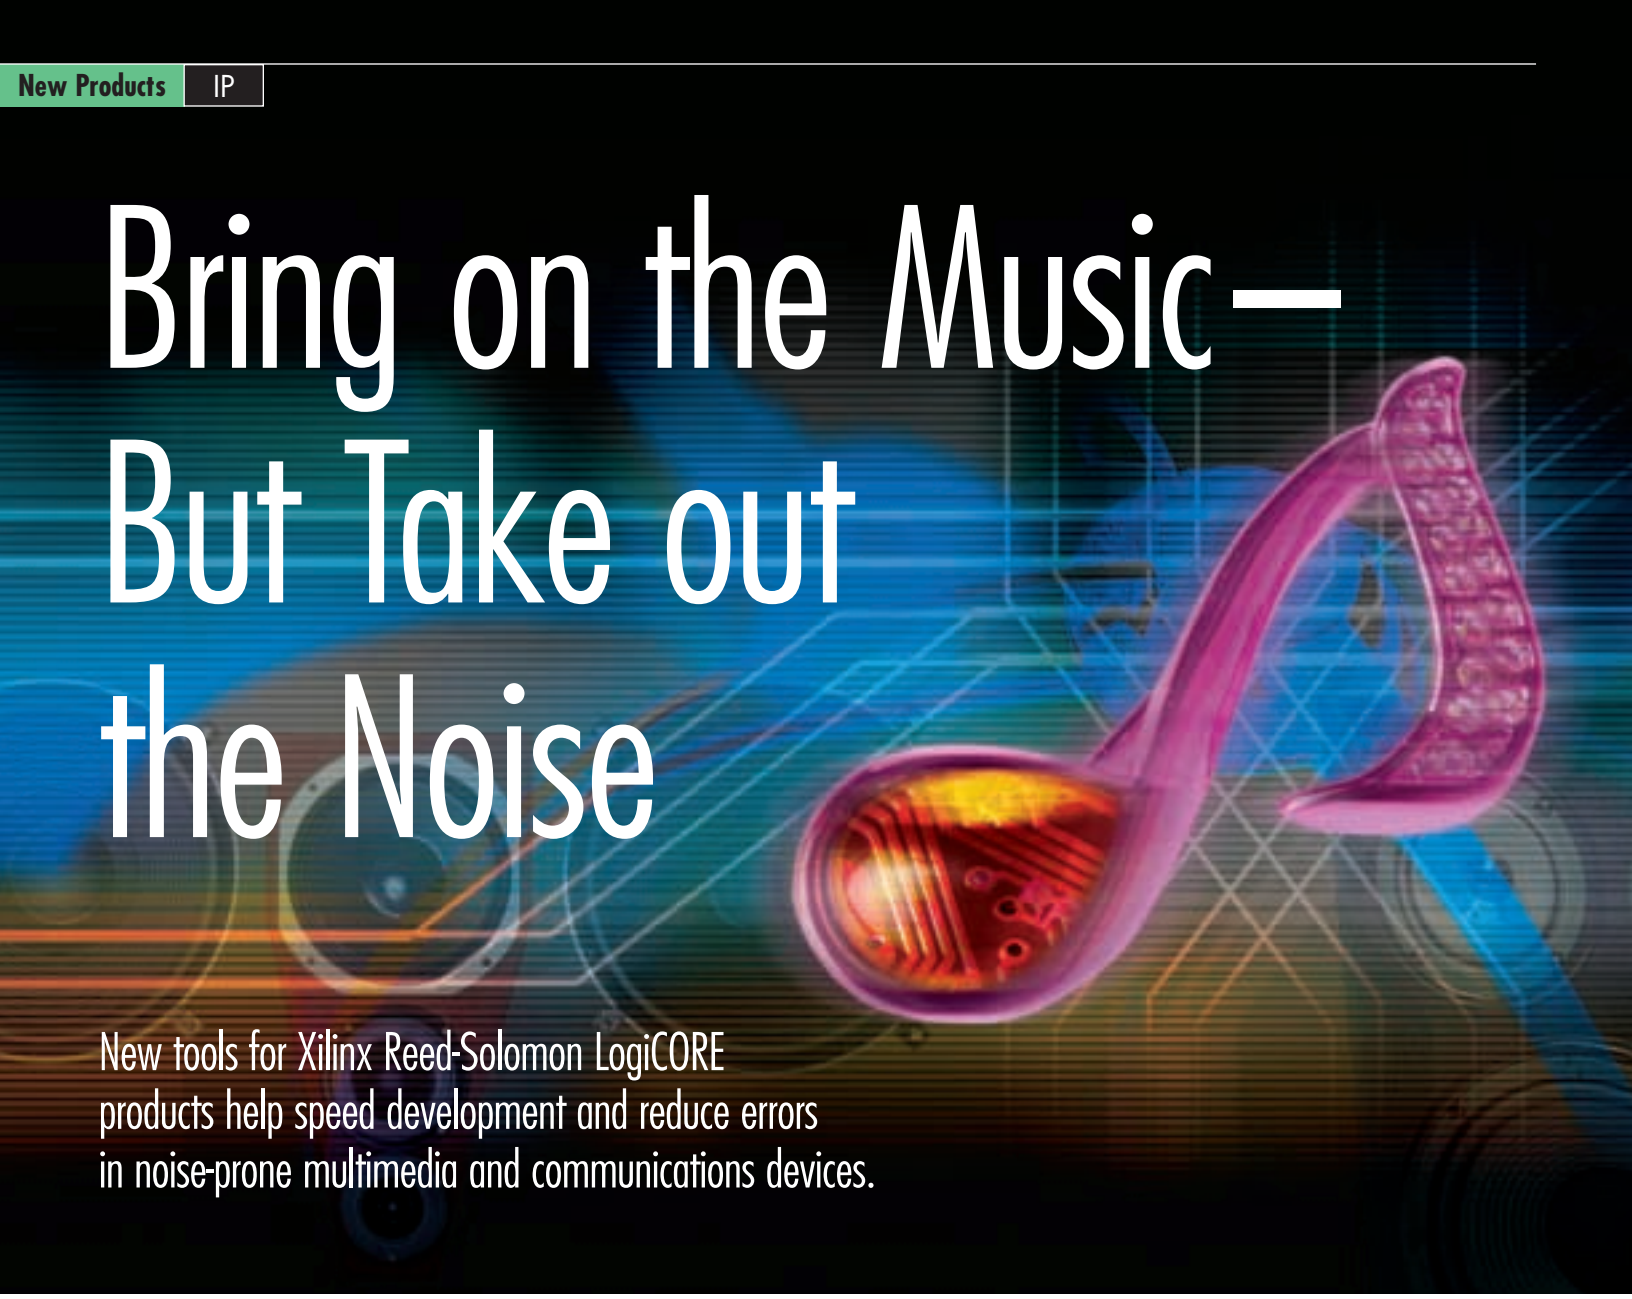

New tools for Xilinx Reed-Solomon LogiCORE products help speed development and reduce errors in noise-prone multimedia and communications devices.

by Warren Miller
VP of Marketing
Avnet Design Services
warren.miller@avnet.com

Because Reed-Solomon codes are very good at correcting burst errors, they are used in applications where noise or defects can take out a series of information bits. Products like CD and DVD players, as well as wireless and hard-wired communications systems, use Reed-Solomon codes.

Successful communications in noisy environments using Reed-Solomon codes are possible through their means of detecting and correcting multiple errors without needing to retransmit the data. This increases bandwidth and improves data integrity in error-inducing communications channels.

Avnet Design Services has created a hardware reference design to demonstrate the functionality of two Xilinx LogiCORE™ Reed-Solomon IP cores in an error-prone system. The reference design uses standard Avnet-developed evaluation kits with the Raptor/ISD (in-system developer) IP environment from Experience First Inc.

Strength from the Core

The strength of Reed-Solomon codes in handling multiple bit errors makes them one of the most efficient and common error correction codes used in communications applications. Other codes (such as Hamming codes) are better suited for the more random error patterns found in memory systems, where only one bit at a time is affected.

Reed-Solomon codes append a series of symbols (check words) to a series of information symbols (data words). The check words are constructed (similar to the familiar parity check) in such a way that if errors are introduced anywhere in the data, the correct data can be reconstructed. The number of errors that can be corrected is one half the number of check symbols (rounded down). For 8-bit-wide data words, a common code has a block size of 207 words with 20 check symbols and 187 data symbols. This is the example code we will use in the reference design.

The Avnet Virtex Development Kit

The Virtex™ Development Kit developed by Avnet includes hardware tools on a single board for testing and debugging a refer-

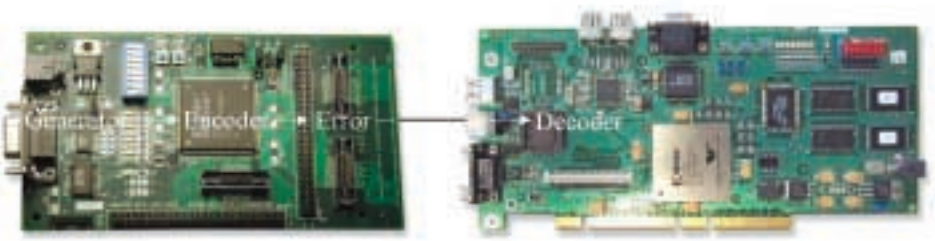

Figure 1 - The Avnet reference design hardware for Reed-Solomon code applications uses a Xilinx Spartan-II FPGA Encoder (left) and a Virtex-II Platform FPGA Decoder (right).

ence design implementing Reed Solomon codes on Xilinx Virtex and Virtex-II Platform FPGAs. The Avnet board includes the following features:

- Data Generator, which creates an input data stream
- Reed-Solomon Encoder, which takes blocks of input data and creates the check symbols
- Error Generator, which injects controlled pseudo-random error patterns into the communications channel
- Reed-Solomon Decoder, which detects and corrects errors in the transmitted data.

The Data Generator and Encoder reside on the Avnet-designed Xilinx Spartan™-II FPGA evaluation board, as shown in Figure 1 (left). The data block with appended check words is transferred to the Virtex evaluation board where the data is received and corrected by the Decoder, as shown in Figure 1 (right).

The Error Generator uses a pseudo-random number generator to select errors to inject into the information block. Adding errors to a reference design is necessary, because when there are no errors, it is not clear that anything is actually being done by the design. An IP development tool that works with the Avnet Virtex development board allows you to easily see the results of the decoder's processing in the reference design.

Raptor IP Development Environment

The Raptor IP Development Environment is an integrated software and firmware tool that works in conjunction with the Avnet Virtex Development Kit. The Raptor environment accelerates the development, integration, and test of IP or IP-based Xilinx FPGA designs. Raptor automatically inserts logic around any IP core to route inputs and outputs to real-time input signals, output signals, and to buffer memory available on the development board. This allows the core to be exercised at "hardware speeds."

Core output signals that are stored in the buffer memory can be read over the USB port to a host computer and displayed on the waveform viewer, just like software simulation results. This hardware speed approach to verification reduces the design/debug cycle time dramatically, making it easy to change the design or test set-up and see the results immediately. Used in conjunction with hardware-based input stimuli, verification of very complex and robust test suites are orders of magnitude faster than approaches based on pure software simulation.

The Raptor user interface, shown in Figure 2, allows you to specify the inputs and outputs to your IP. The necessary logic is automatically added around your IP and then compiled using the Xilinx development tools (Foundation™ Series software in this example). The IP core is then exercised by the hardware and the results are read out of memory and transferred to the PC over the USB port. The signals can be displayed on a waveform display (a sample output is shown in Figure 3). Because the signals are run at hardware speeds, hundreds of thousands of cycles can be exercised in the same time as a single cycle of software-only simulation.

Conclusion

Reed-Solomon codes are a basic building block of many digital communications systems. The Xilinx LogiCORE Encoder and Decoder were easily ported to the Avnet development boards, using the Raptor IP development system from Experience First. For more information on Avnet Development Boards visit www.ads.avnet.com or contact Warren Miller at Avnet Design Services (*warren.miller@avnet.com*). For more information on the Raptor ISD visit expfirst.com or contact Bob Read at Experience First (*bread@expfirst.com*).

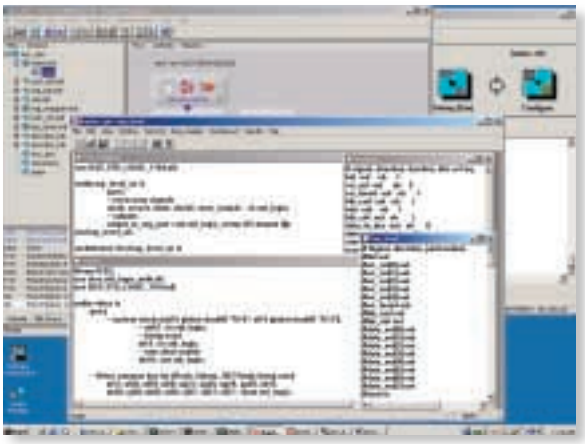

Figure 2 - Experience First's Raptor user interface lets you specify input and output signals to your IP reference design.

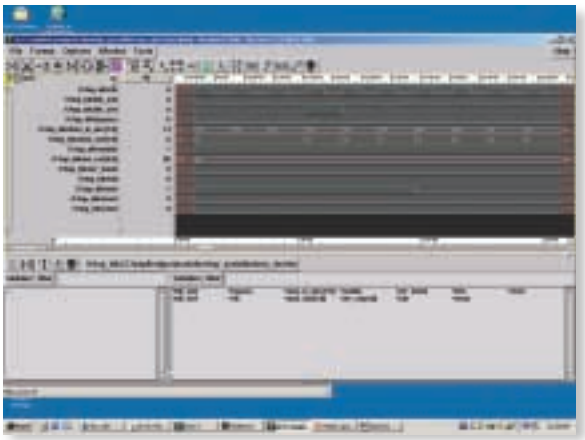

Figure 3 - Real-time Decoder waveform signals are input through the USB port.

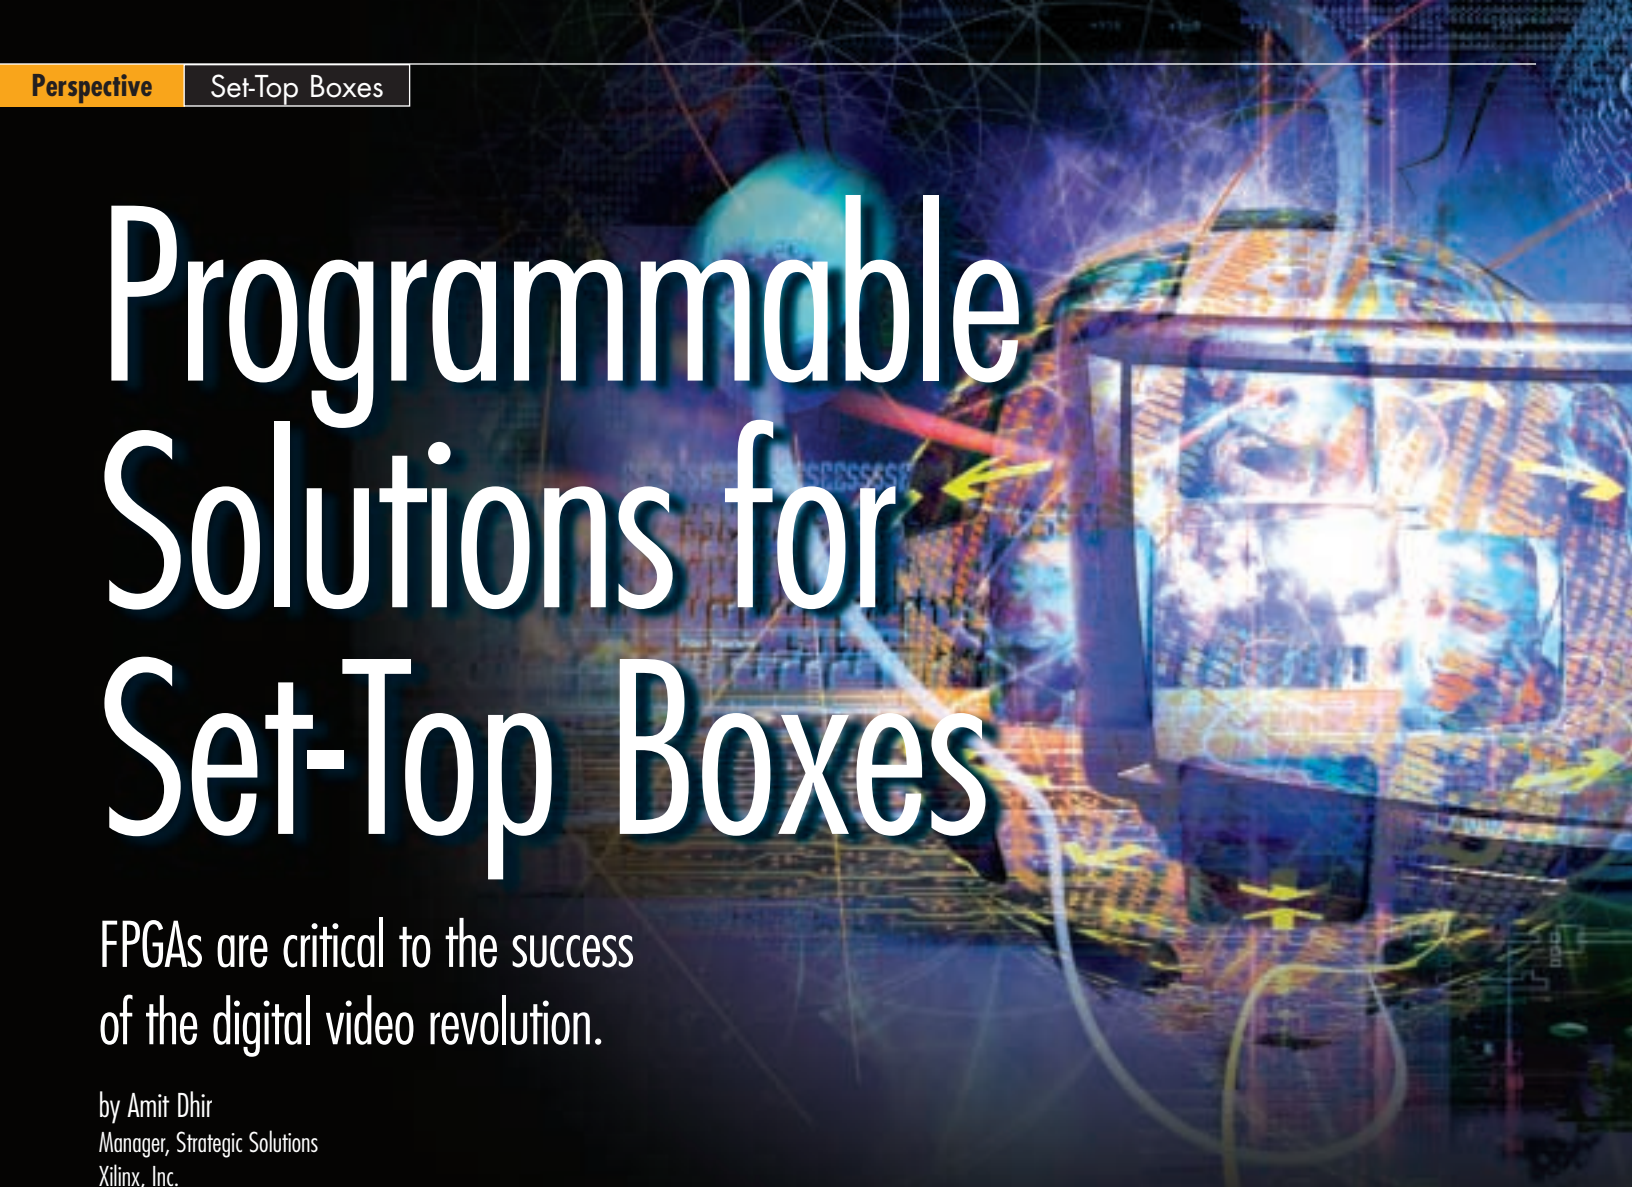

Programmable Solutions for Set-Top Boxes

FPGAs are critical to the success of the digital video revolution.

by Amit Dhir
Manager, Strategic Solutions
Xilinx, Inc.
amit.dhir@xilinx.com

In the early 1970s, the only piece of equipment needed for watching TV was a standard television. In the 1980s, this simple model began to change. Cable and satellite TV providers required consumers to connect their TVs to dedicated networks. Also, operators decided to scramble TV signals, requiring a special box to de-scramble the signals at the consumers home. Today, digital television requires a set-top box to receive and decode digital transmissions.

The main function of a set-top box is to receive additional digital transmissions (cable, satellite and/or terrestrial channels) and to decode into a form suitable for display on analog television sets. It's a complex electronics device comprised of a myriad of hardware and software components, usually connected to the TV set and the cable connection on the wall. They are installed and configured by the local cable, terrestrial or satellite service provider.

The home audio-video landscape is quickly transitioning from analog to digital, now providing several hundred channels of 24-hour coverage, Internet access, and other services. The convergence of functions provided by the television and the PC requires a platform to provide these services. With the arrival of several services, the set-top box is growing beyond its traditional function of enabling digital video. Future generations of set-top boxes will provide more services such as the ability to pause, record, store and replay live video; provide video on demand (VoD); provide Internet access; and control other consumer devices. Providing these services will require the relatively simple set-top box to add components such as flash memory, hard-disk drives, security chips, home networking chipsets, and so on. Programmable logic solutions will enable the integration of these services and components in future generations of set-top boxes.

The Set-Top Box Today

Digital TV set-top boxes are sometimes called receivers. A set-top box is necessary to television viewers who wish to use their current analog television sets to receive digital broadcasts. Typical functions implemented in a set-top box include:

- Decoding the incoming digital signal
- Verifying security levels as well as content access rights
- Separating the audio and video data from the decoded signal
- Decoding the separated audio and video data
- Presenting video to the display device
- Presenting audio to the audio outputs
- Processing and rendering Internet content and other interactive services
- Providing electronic program guide and remote control features.

Figure 1 shows the block diagram of a generic current generation set-top box.

According to Dataquest, sales of digital set-top boxes reached 25.1 million units in 2000, worth an estimated $6.3 billion, and will exceed 92 million units and revenues of $11 billion by 2005. It is estimated that 35 million U.S. homes will use digital set-top boxes by the end of 2006, the estimated year ending the transition to DTV.

Digital set-top boxes are used for satellite, cable, and terrestrial digital TV services. They are especially important for terrestrial services because they guarantee viewers free television broadcasting. A set-top box price ranges from $100 for basic features to over $1,000 for a more sophisticated box. It is often leased as part of signing up for a service. Leading set-top box manufacturers are looking to increase revenues by winning consumers that would like to use the TV set not only to watch television programs but also to browse websites, send e-mail, and shop for goods and services through the Internet.

Set-top boxes in the future will provide more channels and increased choices through specialist channels, which can provide immediate feedback to broadcasters. The set-top box will provide new services such as improved pay-as-you view, online shopping, interactive TV, video-on-demand, and hard-drive storage.

Providing Digital Video Processing

In the digital TV realm, a typical digital set-top box contains one or more microprocessors for running the operating system (possibly Linux or Windows), and for parsing the MPEG transport stream. A set-top box also includes RAM, an MPEG decoder chip, and more chips for audio decoding and processing.

The contents of a set-top box depend on the digital TV standard used. European DVB-compliant set-top boxes contain parts to decode COFDM transmissions while ATSC-compliant set-top boxes contain parts to decode VSB transmissions. The set-top box provides improved quality, support for HDTV transmission, prevents "ghosting or interference effects, and allows broadcasters to provide flexibility in terms of bandwidth vs. quality.

The MPEG-2 encoder, at the heart of every set-top box, is composed of a number of discrete algorithmic sections:

- **Temporal processing** – It seeks out and removes temporal redundancy. This involves storing several successive images and performing motion estimation, compensation, and simple algorithmic processing to derive a pixel-by-pixel difference signal.
- **Spatial Processing** – It uses DCT (Discrete Cosine Transform) to remove the high frequencies not discernable by the human eye.

Statistical or variable length encoding (VLC) is used to remove redundancy in the output from the DCT. The MPEG-2 algorithm makes use of the DCT/IDCT algorithm. DCT returns the discrete cosine transform of "video/audio input" (referred to as the even part of the Fourier series) and converts an image or audio block into its equivalent frequency coefficients. IDCT (or inverse DCT) reconstructs a sequence from its DCT coefficients. Compression allows increased throughput through the transmission medium. Video and audio compression makes multimedia systems very efficient by increasing CPU bandwidth, providing higher video frame rates and better audio quality, and enabling multimedia interactivity.

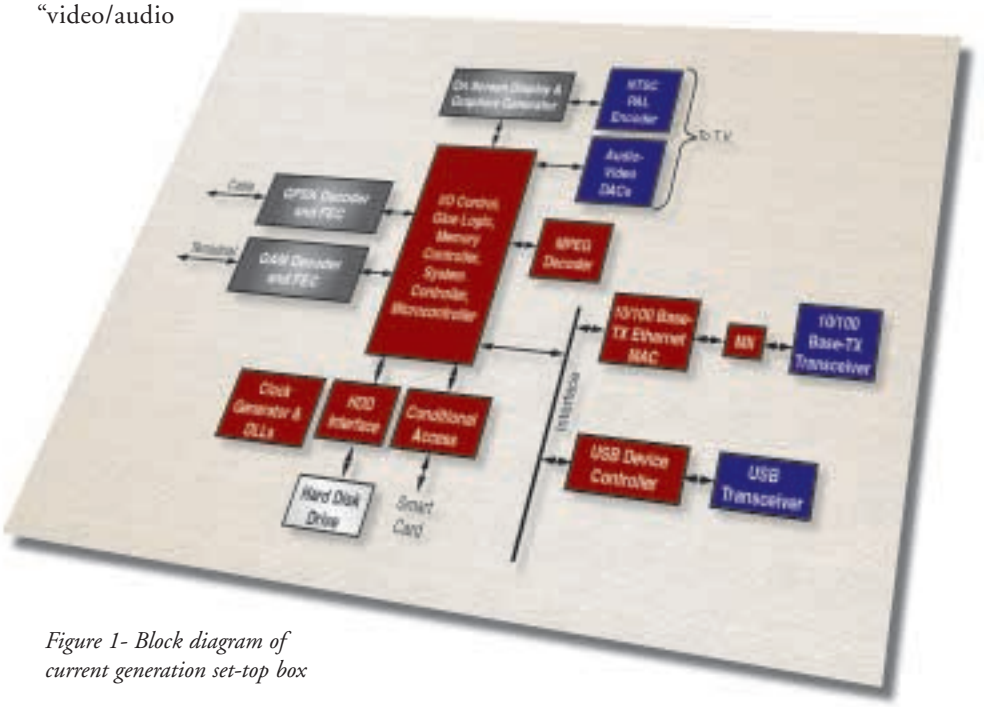

Figure 1- Block diagram of current generation set-top box

While MPEG-2 is the most common compression scheme used in set-top boxes, MPEG-4 compression technology is gathering acceptance. It is earning keen interest from set-top-box vendors and semiconductor companies hungry to add features (such as picture-in-picture) to current designs, and from service providers eyeing it for home networking and for set-tops integrated with digital video recorders (DVRs).

The MPEG-4 initiative is led by satellite providers exploring ways to reduce the bit rates of their broadcast streams and by cable operators looking to add object-based interactive features to their programming. Using the MPEG-4 Advanced Coding Efficiency profile, satellite broadcast streams – currently delivered at bit rates of 2.5 to 3 Mbps – can be trimmed to 1.6 Mbps while maintaining good-quality video. Meanwhile, some service operators are also considering MPEG-4 as a way of saving hard-disk-drive space in set-top boxes equipped with DVRs.

Others see MPEG-4 as effective in reducing the bandwidth required for real-time video for home-networking applications.

There is a need for MPEG-4 in picture-in-picture applications – while a national game broadcast occupies the screen, for example, a local game compressed in MPEG-4 could be broadcast simultaneously so that it appears as a picture within the picture. Further, Internet-streaming content – compressed in MPEG-4 or Windows Media – could also be displayed on screen in a picture-in-picture format.

Most set-top box manufacturers are not looking to replace MPEG-2 with MPEG-4, but to equip a set-top box to transcode MPEG-2 streams into MPEG-4. Clearly, the versatility of the MPEG-4 standard is playing to its advantage as MPEG-4 finds its way into different set-top uses.

Integrating DVR, VoD, and NetTV Functions

The last few years has seen the introduction of the digital video recorder (DVR), also known as digital VCR and personal video recorder (PVR). The DVR uses local storage (such as a hard-disk drive) to enable the user-controlled storage and playback of live digital video streams. The functionality includes the ability to simultaneously record and playback separate video streams or different portions of the same stream in real-time.

With a built-in modem, the DVR device dials a service provider and downloads the programming guide and other software updates. While high unit growth is being predicted for these DVRs, most set-top box manufacturers are looking to incorporate a hard-disk drive and DVR functionality within the set-top box. This will provide the capability to store real-time TV broadcasts in high digital quality, instant access to the recorded data, proactive and quick TV management, and the simultaneous use of multiple data streams. It will also allow the ability to download software and other applications provided by a digital TV service provider.

Apart from providing real-time TV broadcast recording capability, set-top boxes will provide the consumer with video-on-demand capability. The consumer will be able to order a movie instantly. It provides the ability to pause, fast forward, and rewind the movies as often as the consumer desires.

The NetTV is a TV-centric consumer appliance that provides Internet access while using the TV as the primary display. In its most basic offering a NetTV provides limited interactive electronic programming guides or customized information tickers. Advanced NetTVs provide full graphical Web browsing and video email. While the consumer is very interested in services such as access to the Internet, cost remains a big inhibitor for the success of the NetTV appliance. Set-top box manufacturers are looking to integrate the functionality of the NetTV in next generation set-top boxes. In the Internet realm, a set-top box is really a specialized computer that can "talk to" the Internet – that is, it contains a Web browser (which is really a Hypertext Transfer Protocol (HTTP) client) and the Internet's main program, TCP/IP. The service to which the set-top box is attached may be through a telephone line or through a cable TV company.

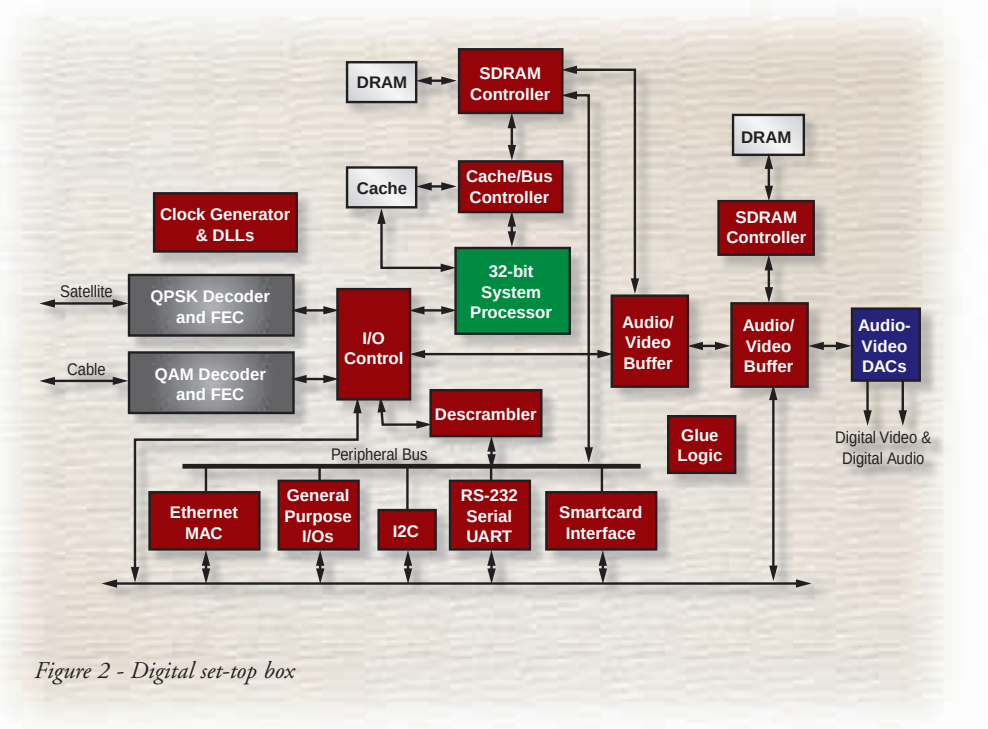

Figure 2 - Digital set-top box

Enabling Broadband Access and Home Networking

The arrival of digital media such as data, voice, video and communications (Internet) through appliances such as digital cameras, MP3 players, cellular phones, and Web pads, is bringing a need for networking consumer devices and PCs. All these consumer devices are demanding high-speed access to the Internet. A house however will have a single broadband access point and the digital media will be shared not only between friends and family across the Internet but also between appliances in the house.

The set-top box of the future will provide high-speed Internet access from satellite, cable, DSL, fixed wireless, and terrestrial access technologies. While one box may support a single access technology, combinations of broadband access technologies will exist within one set-top box. Set-top boxes will also provide interconnectivity between consumer devices through a number of home networking technologies, such as:

- **No new wires** – Phonelines, powerlines
- **New wires** – IEEE 1394, IEEE 1355, Ethernet, USB 1.1/2.0, Optic Fiber, RS-232, IEEE 1284 Parallel Port
- **Wireless** – HomeRF, Infrared, Bluetooth, DECT, IEEE 802.11b, IEEE 802.11a, HiperLAN2

As home networking gains popularity, the set-top box will grow from enabling digital TV and broadband access in the house, to a residential gateway that manages and controls other information appliances in the home.

Categories and Evolution to Home/Residential Gateway

This huge installed base of set-top boxes can be broadly classified into the following categories:

- Analog set-top boxes perform the function of receiving, tuning, and de-scrambling incoming television signals. These receivers have changed very little over the past twenty years.
- Dial-up set-top boxes allow subscribers to access the Internet from the comfort of their living room through the television.
- Entry-level digital set-top boxes are capable of receiving broadcast digital television that is complemented with a pay-per-view system and a very basic navigation tool. They have no return channel, and therefore do not interact with computer servers. Characteristics of this type of low-cost box include limited quantities of memory, interface ports and processing power. Figure 2 shows a digital set-top box.
- Mid-range digital set-top boxes include a return (back) channel, which provides communication with a server located. These set-top boxes provide e-commerce, Internet browsing, and multimedia services. The return channel further allows for customized broadcasts to the local viewing population. These types of boxes have higher processing power and storage capabilities of entry-level boxes.
- Advanced digital set-top boxes are like multimedia desktop computers, containing much higher processing power than other set-top boxes. Enhanced capabilities in conjunction with a high-speed return path can be used to access a variety of advanced services such as video teleconferencing, home networking, IP telephony, VoD, and high-speed Internet TV services. Additionally, it has enhanced graphical capabilities to receive high definition TV signals and can store video on a hard disk drive, while providing the capability to record and view video simultaneously. Such receivers have a range of high-speed interface ports, and resemble residential gateways. For cable, terrestrial, and satellite companies, set-top boxes that support advanced technologies are an opportunity to increase revenue streams through providing services.

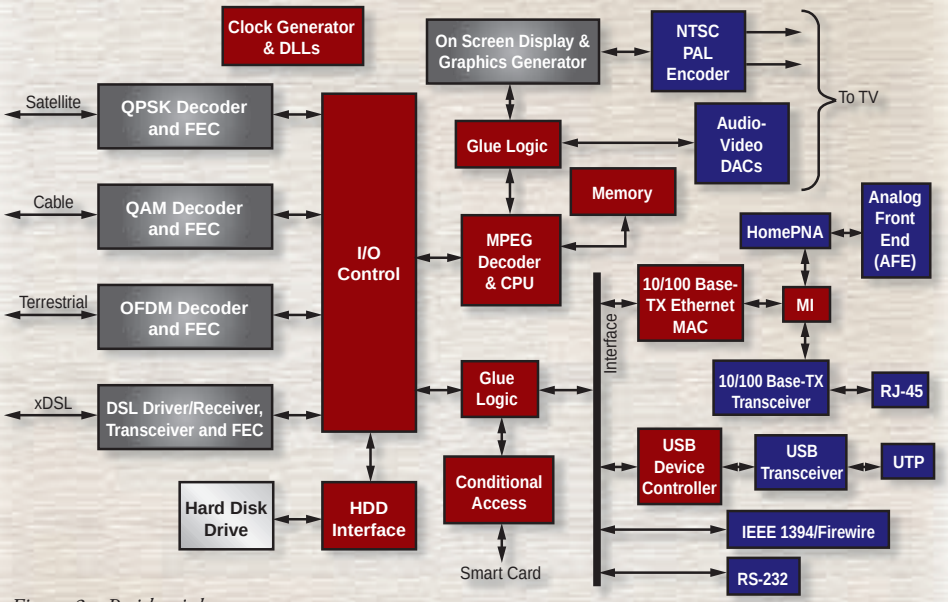

Figure 3 - Residential gateway

The residential gateway is a platform to bring broadband into the house and connect or bridge to home networking technologies. It enables communication between networked appliances in the home and across the Internet. The evolution of new data broadcasting services has created the need for a device to pass digital content between the Internet and the home network. Future gateways will provide integrated services such as remote management, home automation, and home security.

Mass deployment of gateways will come in three distinct phases. While the gateway is a new term, it already exists in many homes. Most of our homes have a set-top box for receiving television and a modem to connect to the Internet. Phase two includes advanced features such as broadband connectivity, home networking interfaces, and IP telephony in the residential gateway. The third phase will be deployment of powerful gateways, capable of delivering video, voice and data throughout the home and providing services such as home automation, energy management, security control, and so on.

Supporting multiple technologies makes the gateway less likely to become obsolete. Support for modularity will fuel the evolution of gateways into a type of application server that consumers will use to distribute broadband services throughout their homes. The gateway must have a reliable and robust hardware platform, and software that is not susceptible to errors. Unlike PC users, consumers will not stand rebooting their gateways. Supporting multiple services with complete security is essential. Functions such as e-commerce transactions, and remote home control and access from authorized service providers are critical. Providing quality of service to support multiple intelligent devices from different vendors is extremely important.

Figure 3 shows the block diagram of a residential gateway. The gateway provides a unified platform to satisfy the needs of most consumers, providing infotainment, conveniences, and communication. The success of the set-top box, and hesitation by a large population to own a PC, allows for the set-top box to grow into a successful residential gateway. Set-top boxes will evolve into multimedia servers, forming the hub of the

home network for primary access to the Internet, such as the residential gateway.

Programmable Logic Solutions Enable Future Set-Top Boxes

For residential gateways to be successful, programmable solutions will have to be at the heart of the system. Programmable logic devices address the fundamental disconnect between ASSPs and provide the interface and protocol translation between different broadband and home-networking chipsets. They provide significant time-to-market advantages and the ability to upgrade quickly, which is imperative for a successful product. This gives you the cutting edge to bring products first to market, while having the ability to remotely add features to the set-top box already deployed in the home.

The set-top box will combine components such as digital modem chipsets, home networking chipsets, processors, memory, and software. Digital modem chipsets provide connectivity to different broadband networks, and home networking chipsets provide interconnectivity technologies between appliances. Other ASSPs/ASICs handle and process digital video and audio services. These ASSPs communicate with each other via buses on the system board. The processor is responsible for coordinating the different components. With additional features, set-top boxes will require higher performance processors to keep pace with increased data throughput. However, all these standards and ASSPs promote differing interfaces making glueless connectivity impossible between the different ASSPs, memories, and processors.

There are three types of set-top box software: operating systems, middleware, and applications. The operating system operates the set-top box parts. The middleware is a layer of software programs that operates between the interactive TV applications and the operating system. Viewers use application software to watch TV and use interactive features. The middleware also enables the smooth interoperation of information appliances and services within the home, to eliminate the complexity, distribution, and technical disparity of the system elements.

For video processing conventional DSPs (digital signal processors) provide a fixed data width and inflexible architecture. They typically have 1-4 MAC units and require serial processing, which limits data throughput. This causes a need for high clock frequency DSPs, which creates system challenges. Hence, multiple DSPs are needed to meet bandwidth requirements, thus causing power and board space issues. Programmable logic devices have a flexible architecture with distributed DSP resources and embedded multipliers. These devices can support any level of parallelism or serial processing through an optimal performance/cost tradeoff. This parallel processing maximizes the data throughput. Hence, programmable logic solutions exceed DSP requirements of the video market – providing both flexibility and performance. FPGAs are off-the-shelf devices, which provide fast time-to-market, rapid adoption of standards, optimal bit widths, and real-time prototyping along with support for high data rates. With a whole suite of DSP algorithms/cores and tools created for programmable logic devices, development time can be reduced by weeks while increasing the performance of the system. For example, the implementation of the DCT/IDCT core in a programmable logic device can off-load the system processor performing MPEG encoding/decoding with a 50X to 200X performance gain.

Programmable solutions also provide the ability to interface to different hard-disk drive types as well as NAND and NOR flash memory types (depending on availability). They also provide content protection capabilities using DES, triple DES, AES, and even proprietary encryption schemes. They also provide system interface functions such as PCI, USB, and so on, in the set-top box and residential gateway. In addition, the presence of the FPGA in the system provides the capability to remotely upgrade features through the Internet when the box has been shipped to the customer, hence providing a significant cost savings.

Summary

The set-top box is driving the digital revolution right into your living room. Your fingertips now command a wealth of high-quality digital information and digital entertainment, right from your favorite armchair. The set-top box revolutionizes home entertainment by providing vibrant television images with crystal clear sound, along with e-mail, Web surfing, and customized information such as stock quotes, weather and traffic updates, on-line shopping, and video-on-demand – right through a traditional television.

The set-top box will evolve into home multimedia centers, possibly forming the hub of the home network system and the primary access for consumers to connect to the Internet. There will be a convergence of technology with equipment connected to the television, such as adding hard drives for television program storage and instant replay. As the set-top box evolves, it will also provide home networking capabilities and value-added services, while becoming the residential gateway of the future.

Programmable logic solutions are necessary for the success of set-top boxes and residential gateways as they provide time-to-market and time-in-market advantages in interfacing disparate technologies, components and system interfaces. They also provide a significant performance advantage over digital video processors. What is clear is that the set-top box market is growing at a very dramatic rate, and when markets are so dynamic and the future is very unpredictable, decreased time-to-market and the ability to upgrade very quickly is imperative for a successful product.

Low-cost Digital Video IEEE 1394 Hardware Reference Design

New reference designs provide OEMs with a complete low-cost solution for integrating digital video in consumer products.

by Xilinx Staff

Xilinx and Convergent Designs recently announced the availability of Centauri, a low-cost digital video (DV) IEEE 1394 hardware reference design based on the low-cost Xilinx Spartan™ FPGA and the Divio NW701 DV codec. The new reference design provides OEM manufacturers with a complete low-cost solution for easily integrating DV over IEEE 1394 into a variety of consumer digital video products such as DVD R/Ws, digital VHS recorders, set-top boxes, residential gateways, Personal Video Recorders (PVRs), Digital TVs (DTVs), video projectors, video editors, and professional digital audio and video equipment.

"The Spartan family of FPGAs provided an ideal low-cost solution," said Michael Schell, president of Convergent Design. "Because of the Xilinx commitment to the Digital Video market, we're able to rapidly introduce cost effective consumer solutions that reduce overall time-to-market and cost."

"The introduction of this hardware reference design is a great example of the compelling value that FPGAs offer the consumer marketplace," said Robert Bielby, senior director of Strategic Solutions Marketing at Xilinx. "Working closely with Convergent Designs and utilizing their digital video and audio design expertise allows us to provide some exciting solutions for the consumer market."

"The Spartan family of FPGAs provided an ideal low-cost solution. Because of the Xilinx commitment to the Digital Video market, we're able to rapidly introduce cost effective consumer solutions that reduce overall time-to-market and cost." — Michael Schell, president of Convergent Design.

Pricing and Availability

Centauri is priced at $450 and immediately available for purchase on the Xilinx eSP website (*www.xilinx.com/esp*).

Xilinx eSP for Digital Video

Xilinx eSP (*www.xilinx.com/esp*) is the industry's first Web portal dedicated to accelerating the design and development of consumer products based upon emerging standards and protocols. Recently updated to contain broad coverage of digital video technologies, the eSP Web portal is a comprehensive resource delivering a powerful array of solutions and information in a single location.

About Convergent Design

Convergent Design has joined the Xilinx XPERTS program, a worldwide third-party Certified Design Service Center trained to take full advantage of the features in Xilinx FPGAs, software, and IP cores.

Convergent Design LLC specializes in the design and development of analog digital video/audio conversion and processing products. Specific technologies include: PCI, 1394, DV/MPEG2 compression, component video, SDI digital video, balanced audio, and AES/EBU digital audio. Convergent Design offers comprehensive design services including initial product specification, schematic capture, FPGA code development, PCB prototype and debug, compatibility testing, and release to contract manufacturer.

Programmable Logic Enables Digital Displays

An overview of an important emerging market.

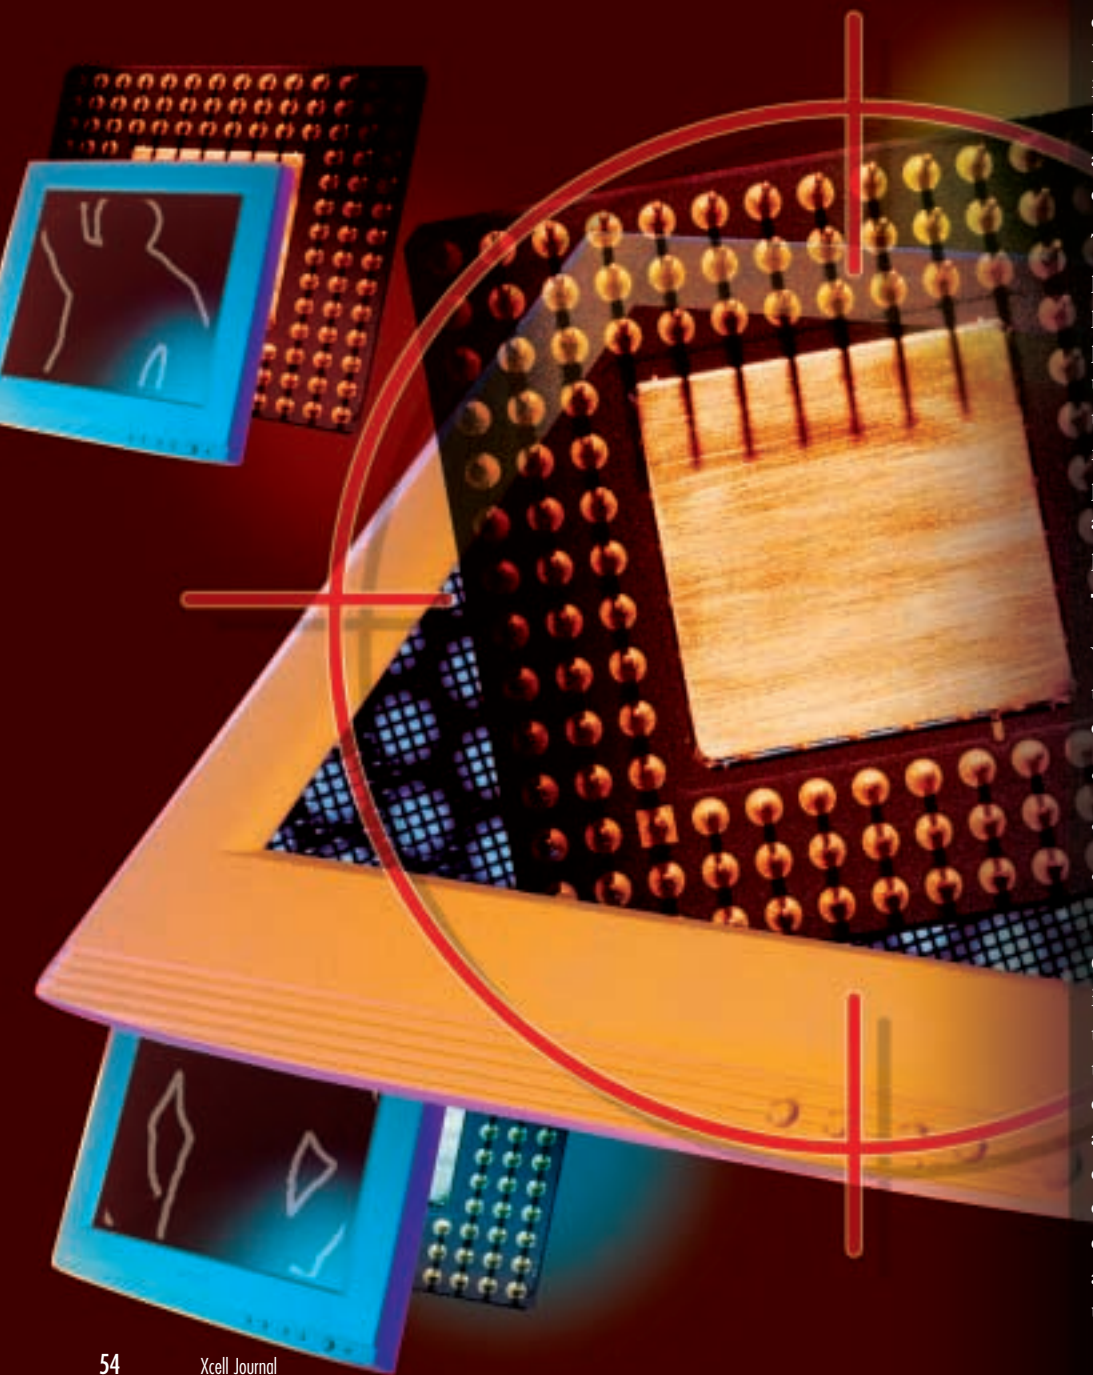

by Mike Nelson
Sr. Manager, Strategic Solutions
Xilinx, Inc.
mike.nelson@xilinx.com

DisplaySearch Inc. forecasts that by 2005, digital displays will eclipse conventional display technologies in market revenue. Digital liquid crystal displays (LCDs), digital plasma display panels (PDPs), digital light processors (DLPs), and many others are fast becoming the display technologies of choice.

This article will explain why this is happening and document the universe of complexity that has been spawned in the process. Next we will examine some of the unique challenges that digital display systems pose for you. Finally, we will review a representative case study to illustrate how programmable logic can be used to your advantage in developing digital display products.

The Digitization of Display Technologies

What's driving the digital display transition? There are three basic answers to this question:

- Content has become digital
- Digital displays achieve superior quality
- Digital display technologies have enabled new and desirable form factors.

Content is king. Analog television became inexpensive because it served a huge content market: broadcast television. But things have changed. The emergence of commodity PCs, the Internet, digital cable and satellite TV, and consumer DVD have combined to establish a huge new universe of digital content. And, as the transition to digital broadcast television unfolds (SDTV and HDTV), virtually the entire display universe will have become natively digital.

With regard to quality, digital content has several advantages.

- Digital content can more effectively filter noise as illustrated by the lack of background hiss on an audio CD as compared to an analog cassette.
- Digital content enables digital manipulation of the content which means that you can effect highly specific enhancements. For example: shadow enhancement, sharpness control, and color manipulations commonly used in medical imaging.
- Digital content can be quickly and repeatedly replicated with precision and widely distributed much more efficiently than analog formats.

And then there is form factor. Without digital displays, products such as today's laptop computer would simply not be possible (remember the first Compaq CRT based portable PCs?). The same applies to wall mounted televisions, PDAs, seat back displays on airplanes, and the mini-displays on the back of digital cameras. The inherent flat panel nature of most digital display technologies is thus a key factor in their increasing success. These applications could not be fulfilled without them.

The Opening of Pandora's Box

While the emergence of digital display technologies has been a good thing, the developments behind them have been the equivalent of opening Pandora's Box. In the frantic rush to digitization, technologies have been invented and re-invented on a multi-dimensional scale across diverse geographies, industries, and suppliers. The result is a multitude of standards, formats, regulations, specifications, and derivatives, all of which makes your job ever more challenging.

To illustrate the complexity that we have wrought, I have compiled a representative list of technologies, formats, and standards introduced as part of the mass digital emergence, in Table 1.

Difficult Questions

While by no means complete, Table 1 nevertheless illustrates a foremost challenge facing you today: risk and complexity. Which are the correct features to support? How do they vary by market segment? How do they vary by geography? How are they evolving? How do you implement the mix of features you need quickly and efficiently? Can you effectively support a family of configurations to service multiple market segments? How do you get to market fast enough to gain market share?

Digital convergence is driving these technologies together, in new ways and in new products – and making that happen is your responsibility.

Unique Challenges of Digital Display Design

Beyond the complexity of their world, environment digital displays also present some unique challenges to you. These include achieving the performance required for the target application, correcting for technology-specific display characteristics, and generating the drive signals for the target display.

Performance is particularly challenging in digital video display applications due to the tremendous computational loads involved.

- You require a high-bandwidth connection from the source.
- You need to perform a complex string of operations on the data. These can include decryption from a secure transmission format, decoding into pixel maps, and optimization of these pixel maps for display.
- You must use the resulting data to generate the driver signals for the display.
- Finally the display driver generates a family of signals that will be distributed to drive each individual pixel. The format of these signals varies with each technology, they have exacting timing requirements, and their specifications are often unique to each and every model.

To make things worse this must all be achieved in real time, and while transiting this pipeline the dataflow will expand from about 25 Mbps (streaming HDTV) to 1.5 Gbps (raw uncompressed 1080i HDTV display data).

The optimization stage of the pipeline is a challenge in and of itself. Here the digital

Wired Connection Technologies	Broadband Access Technologies	Chip to Chip Technologies	LVDS Interconnect Technologies	Television Standards	Streaming Media Formats
IEEE 1394	ISDN	LVTTL	LVDS	NTSC	MPEG-1
USB 2.0	DSL	LVCMOS	LVPECL	PAL	MPEG-2
Ethernet	Cable	Chip to Memory Technologies	BLVDS	SECAM	MPEG-4
HomePNA	WCDMA	HSTL-I	PC Display Formats	SDTV 480i	MJPEG
HomePlug	Chip to	HSTL-III	EGA	SDTV 480p	Real Networks
Wireless Connection Technologies	Backplane Technologies	HSTL-IV	VGA	HDTV 720p	QuickTime
802.11b	5V PCI	SSTL-I	SVGA	HDTV 1080i	MS Media Player
802.11a	3.3V PCI	SSTL-II	XGA	Digital Imaging File Formats	Encryption Standards
HiperLAN2	3.3V PCI-X	SSTL2-I	SXGA	TIFF	DES
Bluetooth	GTL	SSTL2-II	UXGA	JPEG	3DES
HomeRF	GTL+	SSTL3-I	WXGA	Scitex	AES
	AGP	SSTL3-II		Targa	PKI
		CTT		GIF	

Table 1 - Standards and technologies spawned in the mass digital emergence

image data must be adapted to the specific characteristics of the target display technology. This is necessary because while all displays operate on similar principles of color science, each has its own specific (and non-linear) behavioral characteristics. Thus, RGB data (which is most typically targeted for a CRT display) must be processed to display with acceptable results on an LCD, PDP, or other display technology. This processing can be as simple as color correction, or much more involved with algorithms applied for scaling, contrast, brightness, gradation smoothing, edge sharpness, shadow enhancement, and so on. Almost anything is possible, it simply takes adequate processing power.

A Case Study

To illustrate the challenges of digital display design let's analyze a case study example for a digital projector. In the generic case such products traditionally accept analog RGB video input, perform some moderate processing on the data, and then drive the projection display. This is typically effected through a variety of analog (blue), digital (black), and control (green) components as illustrated in Figure 1.

With the advent of digital convergence, the next iteration of such a product may well be required to support some form, or forms, of direct digital input and include the ability to accept and display encoded file formats. In such a case, you face the dilemmas regarding which inputs and formats to support, and then you must select and integrate appropriate components to realize them in a design that meets performance requirements.

Figure 2 illustrates an example for such a design that would support a fast serial USB 2.0 connection as well as an 802.11 wireless LAN connection.

The simplicity of these illustrations belies the complexity of the task. How do you implement the new logic in the system controller to manage the new data flows? What interfaces and signaling standards are required in order to integrate the new components? What extensions to your user interface and control software need to be developed? And, if it is determined that you need to implement and support a variety of these technologies and options, your task becomes much more complex.

The Value of Programmable Logic

Programmable logic is an ideal solution for addressing these challenges. By their nature these devices are flexible – the premium requirement for success in this

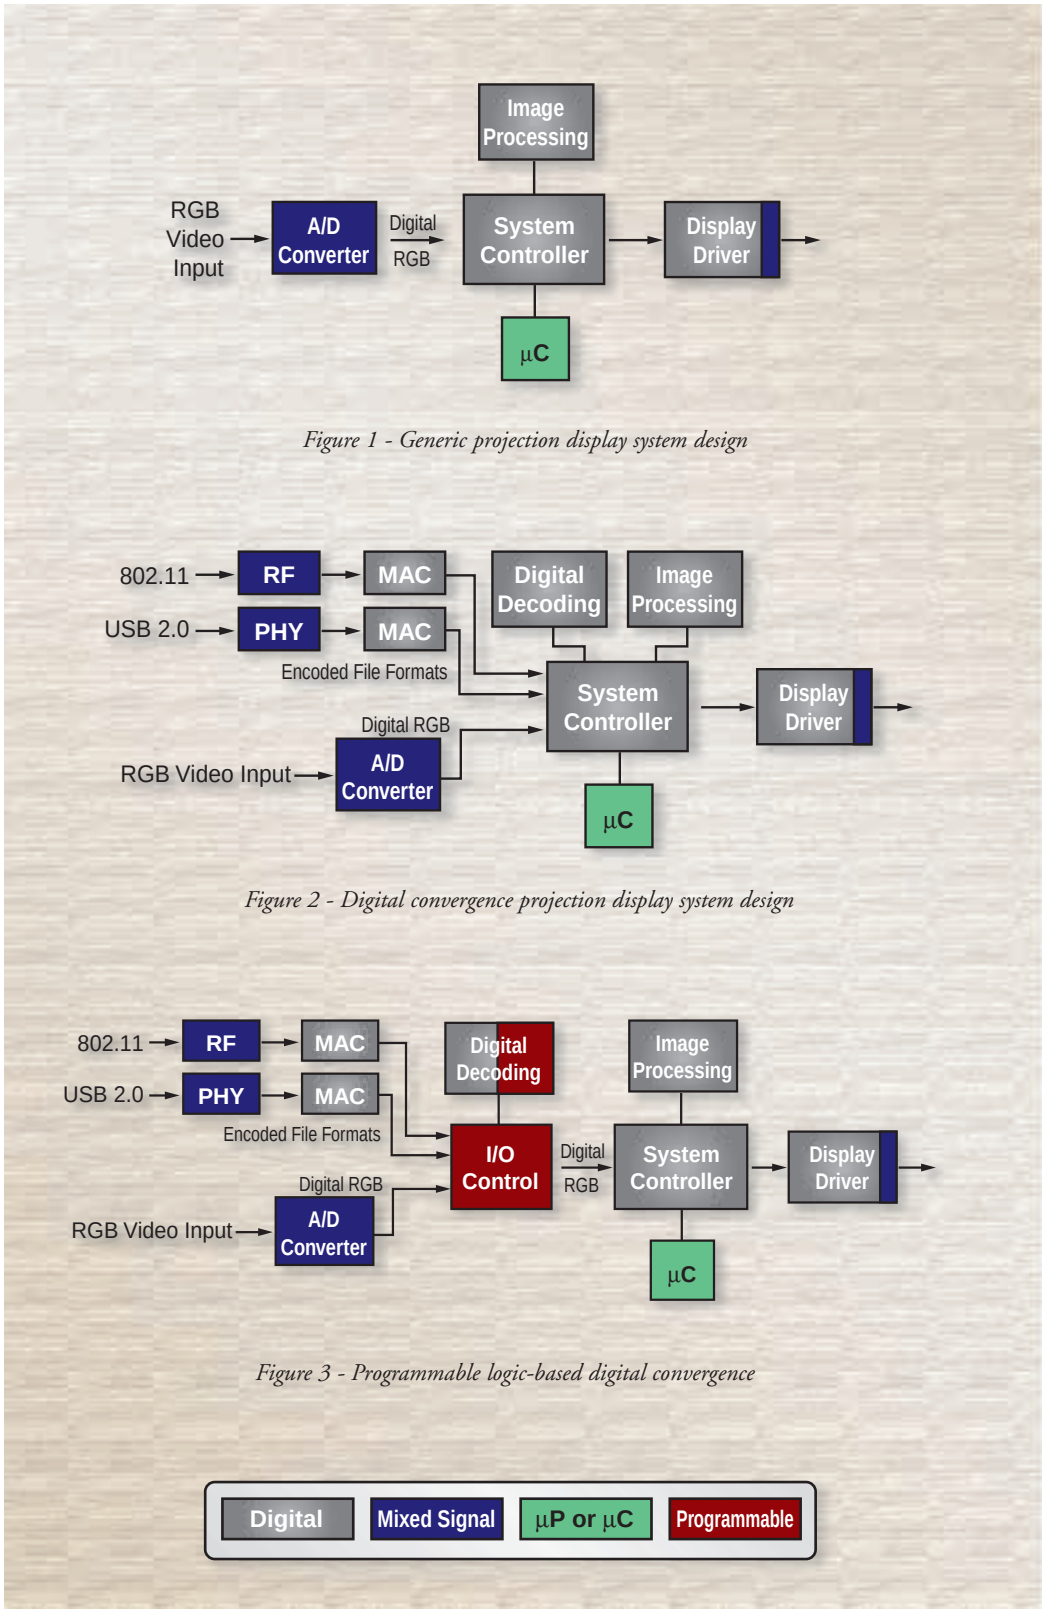

Figure 1 - Generic projection display system design

Figure 2 - Digital convergence projection display system design

Figure 3 - Programmable logic-based digital convergence

endeavor. In addition FPGAs are fast and efficient development platforms, enabling rapid development cycles. Finally, modern FPGAs are extremely cost effective, and therefore viable production solutions for almost any application.

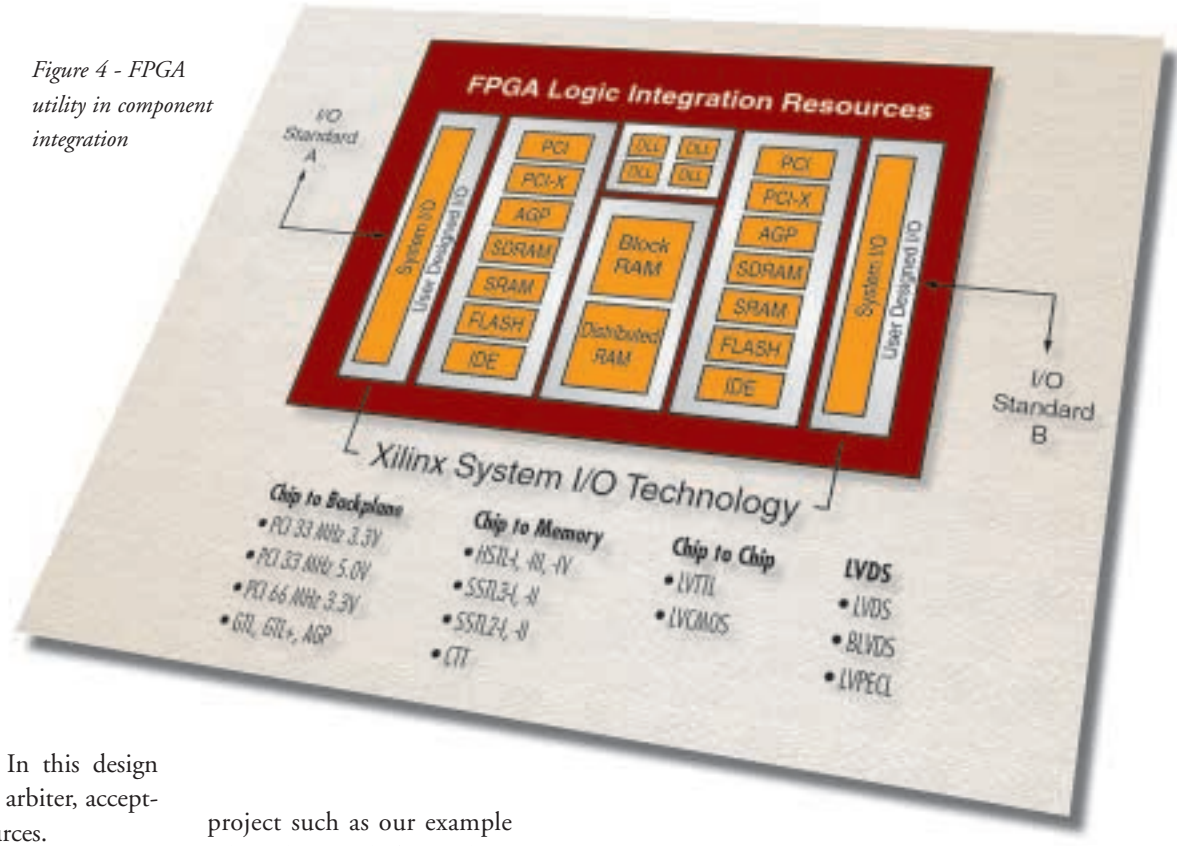

Figure 4 - FPGA utility in component integration

Figure 3 illustrates how an FPGA-based solution could be used to affect our digital convergence projector. As you can see this design inserts an FPGA (illustrated in red) and associated logic between the A/D converter and existing system controller. In this design the FPGA serves as the I/O arbiter, accepting input from all three sources.

In operation the legacy digital RGB signal is simply passed through when this connection is active. In the case of USB 2.0 and 802.11 however, the FPGA serves to manage these new interfaces completely, as well as decode the incoming data stream into the legacy digital RGB format. Decoding can be accomplished entirely in the FPGA or with the assistance of an ASIC or ASSP as appropriate (illustrated by the combined black/red block).

This approach has several important advantages.

- It retains the existing backend of the legacy design essentially unchanged. This bounds development complexity and reduces risk.
- The programmable bridge imposes no schedule penalties for numerous iterations. This can be a significant advantage when you are tasked with integrating new and unfamiliar technologies.
- Upon completion the design can be released and in production very quickly.

How is all this possible? Figure 4 illustrates some of the standard features and IP available in Xilinx FPGAs that make a project such as our example relatively straightforward.

On the perimeter of Figure 4 is System I/O, which allows each and every I/O pin to be programmed to support any of 17 different signaling standards. But System I/O doesn't stop there. In addition to the basic signaling parameters it supports programmable drive strength and multiple slew rates too. These features make it easy to deal with unanticipated PCB behaviors (in those not so rare cases where fabrication reality doesn't match design theory) as illustrated in Figure 5.

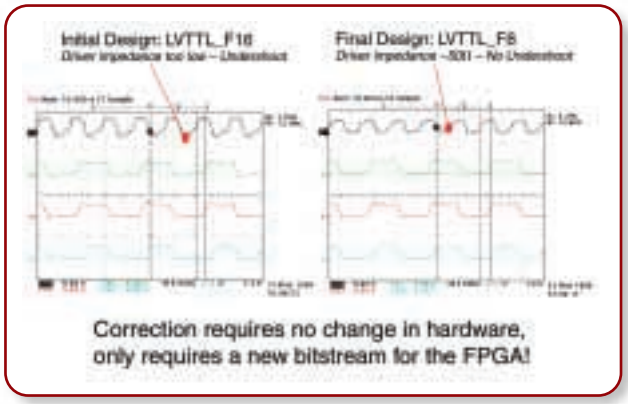

Figure 5 - The benefit of programmable drive strength in System I/O

Some FPGAs, such as the new Spartan™-IIE family from Xilinx, even go a step further including support for LVDS, BLVDS, and LVPECL differential signaling standards at up to 400 Mbps per pin pair. This enables very high-performance component interconnection without the need to resort to higher pin count and more expensive packaging. Further, it reduces system power, lowers EMI, and is much less sensitive to noise as illustrated in Figure 6.

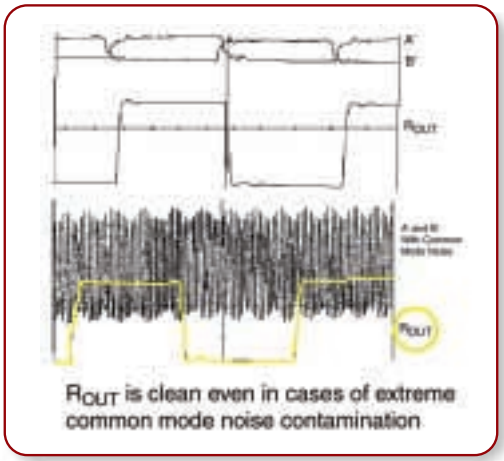

Figure 6 - Noise immunity benefit of LVDS signaling

Figure 4 also shows a representative sample of the standard controller IP modules available for FPGAs. These are commercial quality function blocks that are available to jump-start your design. Much can be accomplished with standard IP as solutions are available to address most of the topics listed in Table 1.

For buffers and FIFOs you have a variety of on-chip memory resources to choose from. These include 200 MHz flip-flops, true dual ported Block RAM, a Shift Register Mode (SRL16) capability in the FPGA's fabric Look-up Table (LUT) structures, and highly configurable Distributed RAM. These features provide high-performance and silicon-efficient solutions for almost any on-chip memory need.

For clock management Xilinx FPGAs feature four or more Delay-Locked Loops (DLL) per device. These provide the resources to synchronize and connect your system elements together and manage EMI. These DLLs exhibit superior noise immunity compared to PLLs, and they feature excellent jitter specifications, making your job easier. A few examples of their utility are illustrated in Figure 7.

Finally, an array of Configurable Logic Blocks (CLB) and internal interconnect resources tie everything together. These are illustrated in Figure 8 and are the underlying fabric that make an FPGA an FPGA.

The FPGA Way

While the benefits of programmable logic are obvious as illustrated in our example, its value can be even greater when leveraged as a foundation element in your design. Figure 9 illustrates how our digital convergence projector might look if it were designed from scratch, only this time the FPGA way.

In this design the flexibility of FPGAs is being leveraged to maximum advantage. By designing the core logic of the system controller into an FPGA component you gain maximum flexibility in the selection of every other component you require – be they HSTL, SSTL, LVTTL, LVDS, or whatever, they can be quickly and efficiently integrated.

Another advantage in our example is the modular architecture for system input. In this design we could support a family of configurations to address a variety of geographic and application requirements. Further, there is no reason why these cannot be developed in a serial fashion, enabling the most important configurations to get to market first. And, all that would change from configuration to configuration in the core design is the bitstream program in the FPGA.

With programmable logic in the data path you have tremendous ability to tailor encoding/decoding, encryption/decryption, and image processing functionality to your precise needs. More importantly, you also have the flexibility to keep up with changes as these needs evolve over time. Take file decryption for streaming media as an example. Today there are no firmly established standards, and the standards that do exist vary widely by geography and content provider. And remember, Content is king, and that using a programmable device as the decryption engine could allow you to support whatever your customers require both today and tomorrow.

When used as the heart of the display driver circuit, programmable logic can give your design the ability to support two or more display options. This can be of tremendous value in managing the cost for this high dollar bill of materials component (more than paying for the FPGA in many cases) or to support a family of products based upon a common core design that increases your accessible market. In addition, you can use LVDS to route these signals around the board (which quite often involves traversing large tracts of real estate) and thus minimize system level EMI and the impact of noise on these critical signals.

FPGAs are also well suited for crafting a unique and attractive user interface for your design. They are the very definition of GPIO (General Purpose I/O) and can implement microcontrollers (or even a PowerPC microprocessor) for supervisory control. In today's competitive markets the user interface can be one of the most effective ways to differentiate your prod-

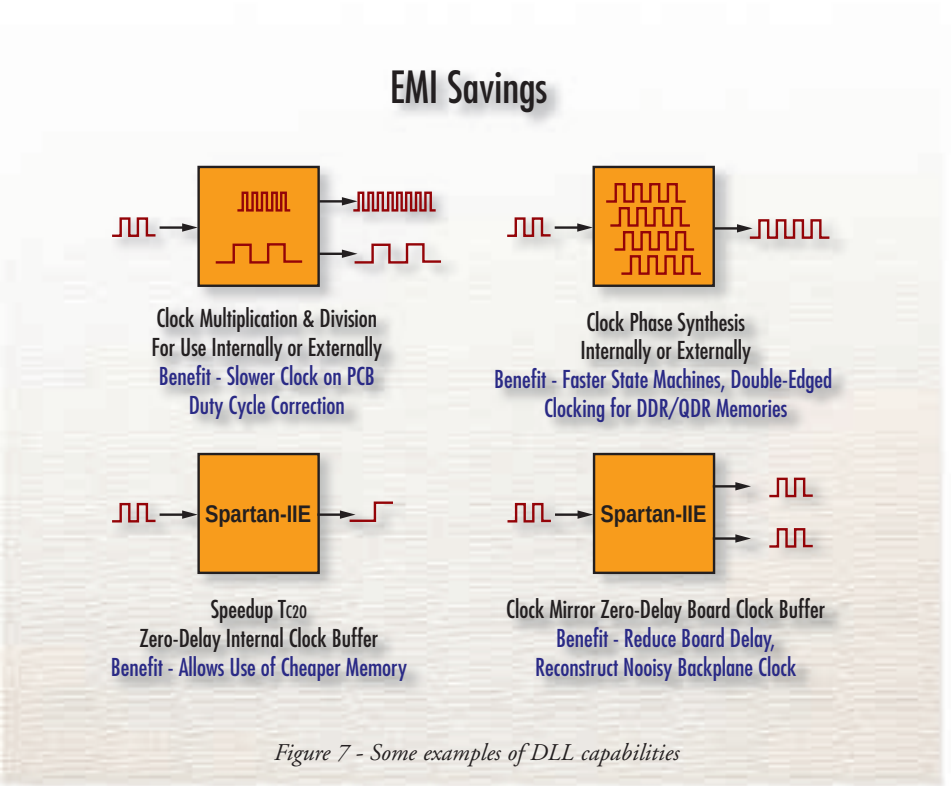

Figure 7 - Some examples of DLL capabilities

uct from that of your competition, and an FPGA gives you the maximum freedom to innovate.

Finally, your FPGA based solution is never frozen. If a customer comes to you requesting a new feature, a slightly different capability, or a new configuration, you have a platform to quickly and efficiently deliver it. When inevitable bugs and incompatibilities crop up you can not only fix them quickly, you can also update deployed systems in the field. This can greatly reduce support and service costs. And if you ever face a supply problem for a system component while in production, you have the flexibility to find and support an alternate solution to keep your factory running, your product shipping, and revenue coming in.

Conclusion

The era of digital convergence is upon us. From pictures to e-mail, from music to news, the world has gone digital. And because of this digital explosion today's systems require ever more connectivity and intelligence. It is no longer good enough to have the best widget or display. Now you need a more digitally connected and data manipulating widget or display, one that supports the standards and formats in your target market and geography – and one that does it before your competition.

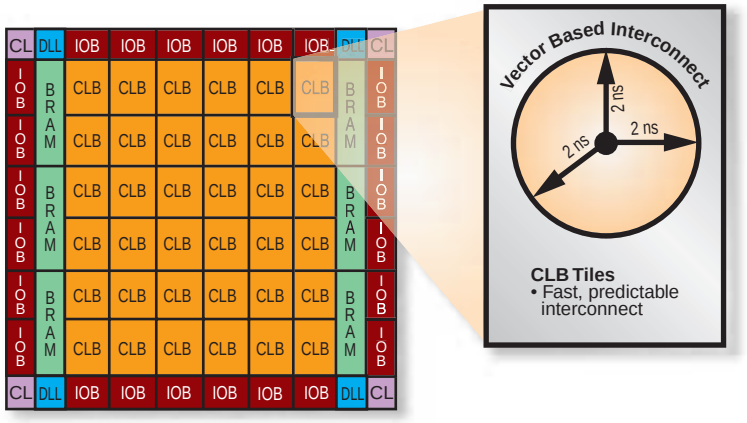

Figure 8 - FPGA CLBs and interconnect resources

Programmable logic is an invaluable asset in confronting this challenge. Its inherent flexibility makes it an ideal mechanism for grafting new functionality into an existing design. Its rich features, efficient development flow, and extensive IP support will simplify your job and give you a chance to meet aggressive schedules. The newest generations of devices are cost efficient solutions for almost any design. And when made a fundamental part of your architecture from the beginning, this technology can be exploited to modularize your configurations, provide flexibility with critical and costly components, and tailor your product functionality to your exact needs.

To learn more about digital video and digital convergence technologies, and how FPGAs can help with them, visit the Emerging Standards and Protocols (eSP) Web portal at *www.xilinx.com/esp*. This website was developed by Xilinx as a resource for the digital system design community and is specifically targeted at dealing with these new and challenging technologies. To date, segments have been deployed for home networking, Bluetooth, and digital video technologies, and more are on the way.

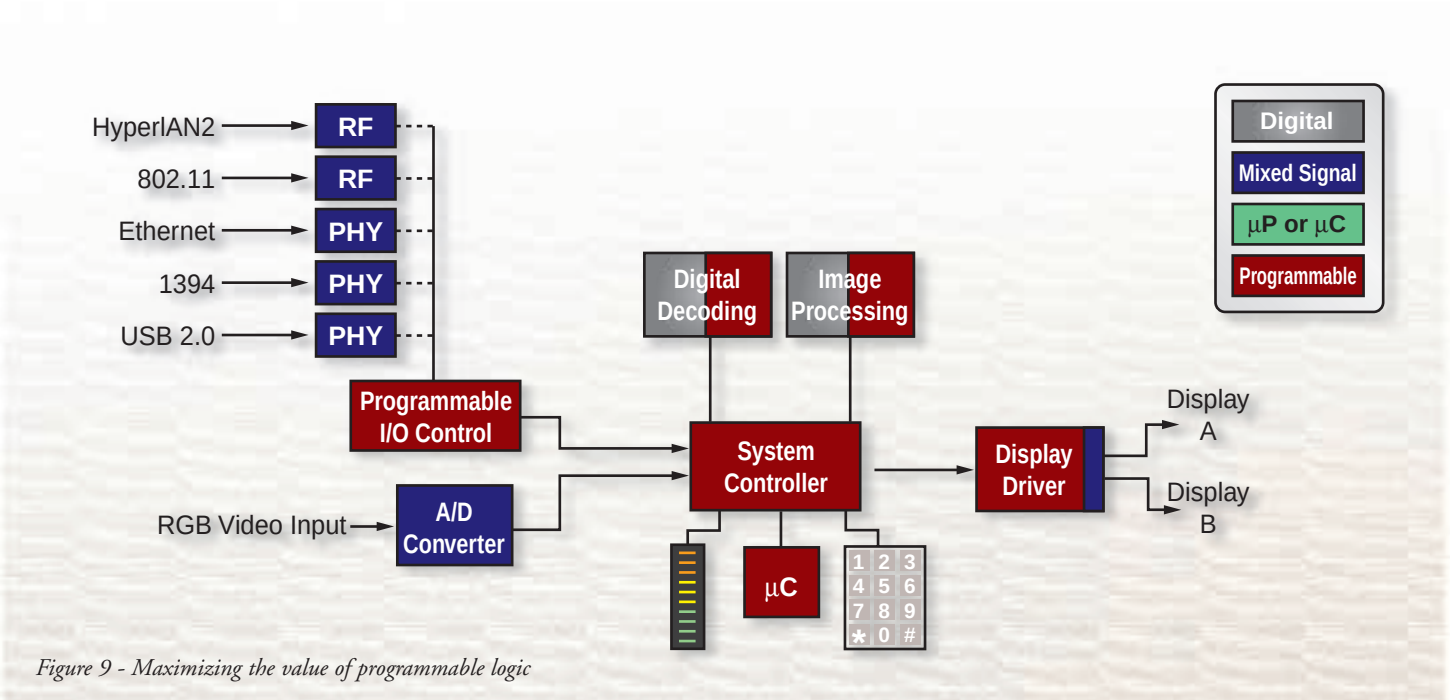

Figure 9 - Maximizing the value of programmable logic

Digital Consumer Convergence Demands Reprogrammability

Not long ago, the idea of creating an ASIC on the desktop was considered a novel concept.

by Robert Bielby
Sr. Director, Strategic Solutions
Xilinx, Inc.
robert.bielby@xilinx.com

Seventeen years ago, Xilinx developed the industry's first Field Programmable Gate Array (FPGA). The concept was quite simple – using Static Random Access Memory (SRAM) as the basis of the logic fabric, you could instantly develop integrated logic circuits directly on your desktop. The benefits were far reaching, because this technology did away with the risk and delay typically associated with designing custom Application Specific Integrated Circuits (ASICs).

While the concept of "user programmable" logic was perhaps first pioneered by enterprising engineers who used Programmable Read Only Memories (PROMs) as a "programmable ASIC", this solution was able to achieve only small levels of integration. The cost efficiency and logic density achieved by using PROMs as custom logic was extremely poor compared to ASICs. Interestingly, however, the concept of using look-up tables to implement logic functions is actually a basic technology used in today's multi-million gate FPGAs.

Today, FPGAs have evolved into a mainstream technology because their current densities and performance compete with custom ASICs. Additionally, designing for an FPGA is almost exactly the same as an ASIC – your circuit can be described either by schematic, or by a high-level hardware description language such as VHDL or Verilog. This description is then synthesized for the FPGA with the final result being a configuration bitstream that ultimately programs device look-up tables that are interconnected by programmable wires. Just like the look-up tables, the configuration bitstream also determines the connectivity for the programmable wires.

Because the programmability is based upon volatile, SRAM technology, there is a requirement to "re-write" the configuration bit stream every time power is removed and re-applied to the FPGA. This configuration bit stream is usually contained in an external, memory which is dedicated for configuration, or by reserving a portion of a main system memory with the FPGA configuration information.

When FPGAs were first introduced, the fact that they required programming from an external memory source, and that they lost their configuration once power was removed, was generally considered a liability. Fixed ASICs had none of these requirements or added complexities. And for the most part, ASICs were still significantly cheaper than FPGAs – at least on a unit cost basis.

However, for applications that shipped a relatively low number of units, it could be shown that the high non-recurring charges for the ASIC more than offset the higher unit costs of the FPGA. Because of these cost constraints, the use of FPGAs was historically relegated to the lower volume, less cost sensitive applications that had large-scale integration requirements.

Over time, the additional benefits of being first to market and the ability to fix bugs or accommodate late specification changes rose in importance, and the decision to select an FPGA over an ASIC became common. This proved to be true even in cases where it could be demonstrated that an ASIC solution provided a lower overall solution cost than the FPGA. However, there were practical limits to how great the unit volumes or disparity in price would become before a design would transition from FPGA to ASIC.

These fundamental benefits of being first to market and of risk aversion have proven to be especially important to the networking and communications industries – as demonstrated by the fact that currently 70% of the $5 billion market for programmable logic is consumed by networking and communications applications. FPGAs have been particularly important in these markets because they have been able to keep up with the pace and innovation in those markets and accommodate the wide range of new standards that continuously emerge.

Flexibility for the Masses

Moore's Law, which says that transistor density of an ASIC will double every 18 months due to advances in semiconductor technology, has played a key role in driving larger density FPGAs with greater features and performance. This has in turn continued to fuel the increased demand and consumption of FPGAs in networking and communications applications.

The advances in semiconductor technology have not only yielded larger and faster FPGAs, they have also yielded FPGAs that are 10,000% cheaper than they were five years ago! The result is that FPGAs, which were previously practical only for prototyping or low volume, high-end applications, are now appearing in some of the hottest, newest high-volume consumer electronics.

Four years ago, Xilinx developed the Spartan series family of FPGAs that were optimized for low-cost, high-volume applications. The results are impressive; from MP3 players, to DVD writers, digital cameras, digital modems, and a host of other consumer electronics – FPGAs have rapidly become a key driver behind the digital consumer revolution.

And reprogrammability, which used to be considered a liability, is now clearly seen as a key benefit in not only getting a product to market sooner, but keeping it longer in the market by providing the ability to upgrade it and add new features once it's in the field. But how and why programmable logic is being used is sometimes just as interesting as where it's being used.

ReplayTV – Personal Video Recorder

Personal video recorders (PVRs) are perhaps one of the most exciting consumer products to offer new features and capabilities made possible only by the combination of digital technologies and FPGAs. With a PVR, traditional analog video programming is converted to digital using MPEG 2 encoding and stored directly to

a hard drive – enabling instant search access and high quality video imaging.

PVRs represent a quantum leap in capability and quality compared to traditional videocassette recorders (VCRs). For example, a PVR such as ReplayTV's can store as much as 60 hours of programming, allowing viewers to watch programming on their personal schedules instead of those set by broadcasters.

The ReplayTV also contains an integrated 56K-baud modem, which is used to download the equivalent of a TV guide that can be searched, sorted, and grouped like traditional database programs. This allows for easy recording setup and unique programming search capabilities. This modem connection also enables additional unique capabilities – such as reconfig-

uring the FPGA. The reprogrammable logic is updated simply by downloading a new bitstream. If the unit is already installed in a customer's home, the bitstream is downloaded from ReplayTV's Internet server. This allows bugs to be fixed even after the customer takes the unit home, and lets ReplayTV add new features as necessary. Obviously, this extends the life of the product too, because rather than having to replace it as market requirements change; the customer simply has the unit's logic reprogrammed via the modem.

ReplayTV reprogramming takes place in the background: each evening the PVR automatically downloads TV schedule information from the company server, along with any bug fixes, operating system updates, or program modifications. The customer's unit is improved as he sleeps without intervention on his part. End users are typically never made aware of changes or fixes to their systems unless the functionality of the PVR changes as new features are added.

In one case, ReplayTV found itself forced to deal with a condition that caused degraded video quality in a few systems already in homes – one of the chips in the system had an undocumented clock threshold switching problem that varied as a function of lot processing. Because the control signals for this device were generated in the FPGA, it was possible to eliminate the problem by changing the timing of the FPGA-generated signal.

The company responded with a software change that was uploaded to all ReplayTV systems in the field as soon as it was debugged and certified. Most customers never realized that the change had been made; though some may have noticed improved video quality. A revolutionary capability enabled by FPGA programmable technology!

KB Gear – JAMCAM 3.0 Digital Camera

Perhaps one of the least likely places one would expect to find an FPGA would be in a toy digital camera. But time to market pressures and the risk of missing the narrow Christmas window of opportunity drove KB Gear Interactive, a manufacturer of Internet communications products and interactive gear for young computer trekkers, to chose the low cost Spartan™ FPGA because of its affordability and design flexibility.

As a testament to the viability of FPGAs in such a cost sensitive consumer product – the JAMCAM 3.0, which cost just $99, was sold in over 14,000 retail outlets including Best Buy, Target, K-Mart, WalMart, Circuit City, and a host of other highly recognizable consumer retail outlets. Because KB Gear had designed a toy that had broad appeal and used a Spartan FPGA, they were able to deliver one of the "must have" toys for the 2000 Christmas season on the shelves in time for the Christmas season. Years ago, the thoughts of an FPGA in a $99 kid's toy would have seemed impossible.

Conclusion

Seventeen years ago the idea of creating an ASIC on the desktop was considered a novel concept; they were primarily limited to prototyping and low volume production. The fact that the FPGAs needed to be programmed every time they were powered up was considered both a liability and a risk. Today, because of their lower cost and high flexibility, FPGAs have found a home in a wide range of applications. Xilinx FPGAs are now used in a wide range of high-volume, cost sensitive applications, especially those that require reprogrammability to meet the continuously changing standards and demands of the new digital consumer markets.

For High-Speed Design, You're Better Connected With Xilinx.

SystemIO Through our XPERTS program, Xilinx certifies many 3rd party design centers specializing in terabit applications. By taking advantage of industry-unique features such as the SelectI/O™-Ultra technology in the Virtex®-II Platform FPGA, Xilinx and its XPERTS partners provide the fastest and most flexible high-bandwith solution supporting rapidly evolving connectivity standards.

Our XPERTS partners have extensive experience using the popular Xilinx PCI 32/33, 64/66 and PCI-X cores, and are ready to take on the challenges of designing with multi-gigabit serial interfaces using the 3.125 Gbps serial I/O technology in the next generation Virtex-II family.

SystemIO addresses all aspects of system connectivity

Xilinx provides the most comprehensive SystemIO solution to address your interface needs ranging from "inside-the-box" to wide area network (WAN) applications.

Using Virtex-II 840 Mbps LVDS performance and an abundance of memory and logic resources, Xilinx is able to provide solutions for 10GE MAC, PCI & PCI-X, RapidIO, POS-PHY Level 3 and 4, Flexbus 4, HyperTransport, and other source-synchronous bus standards.

With SystemIO and our XPERTS program, you're better connected with Xilinx.

For more information visit **http://www.xilinx.com/xperts/systemio.htm**

XILINX®
The Programmable Logic Company℠

www.xilinx.com

FORTUNE® 2002 100 BEST COMPANIES TO WORK FOR

*Flex*Bench Tool Suite Relies on Xilinx Silicon and Software

Improve your time to market with rapid prototyping and system verification enabled by Virtex-II FPGAs and ChipScope Integrated Logic Analyzer.

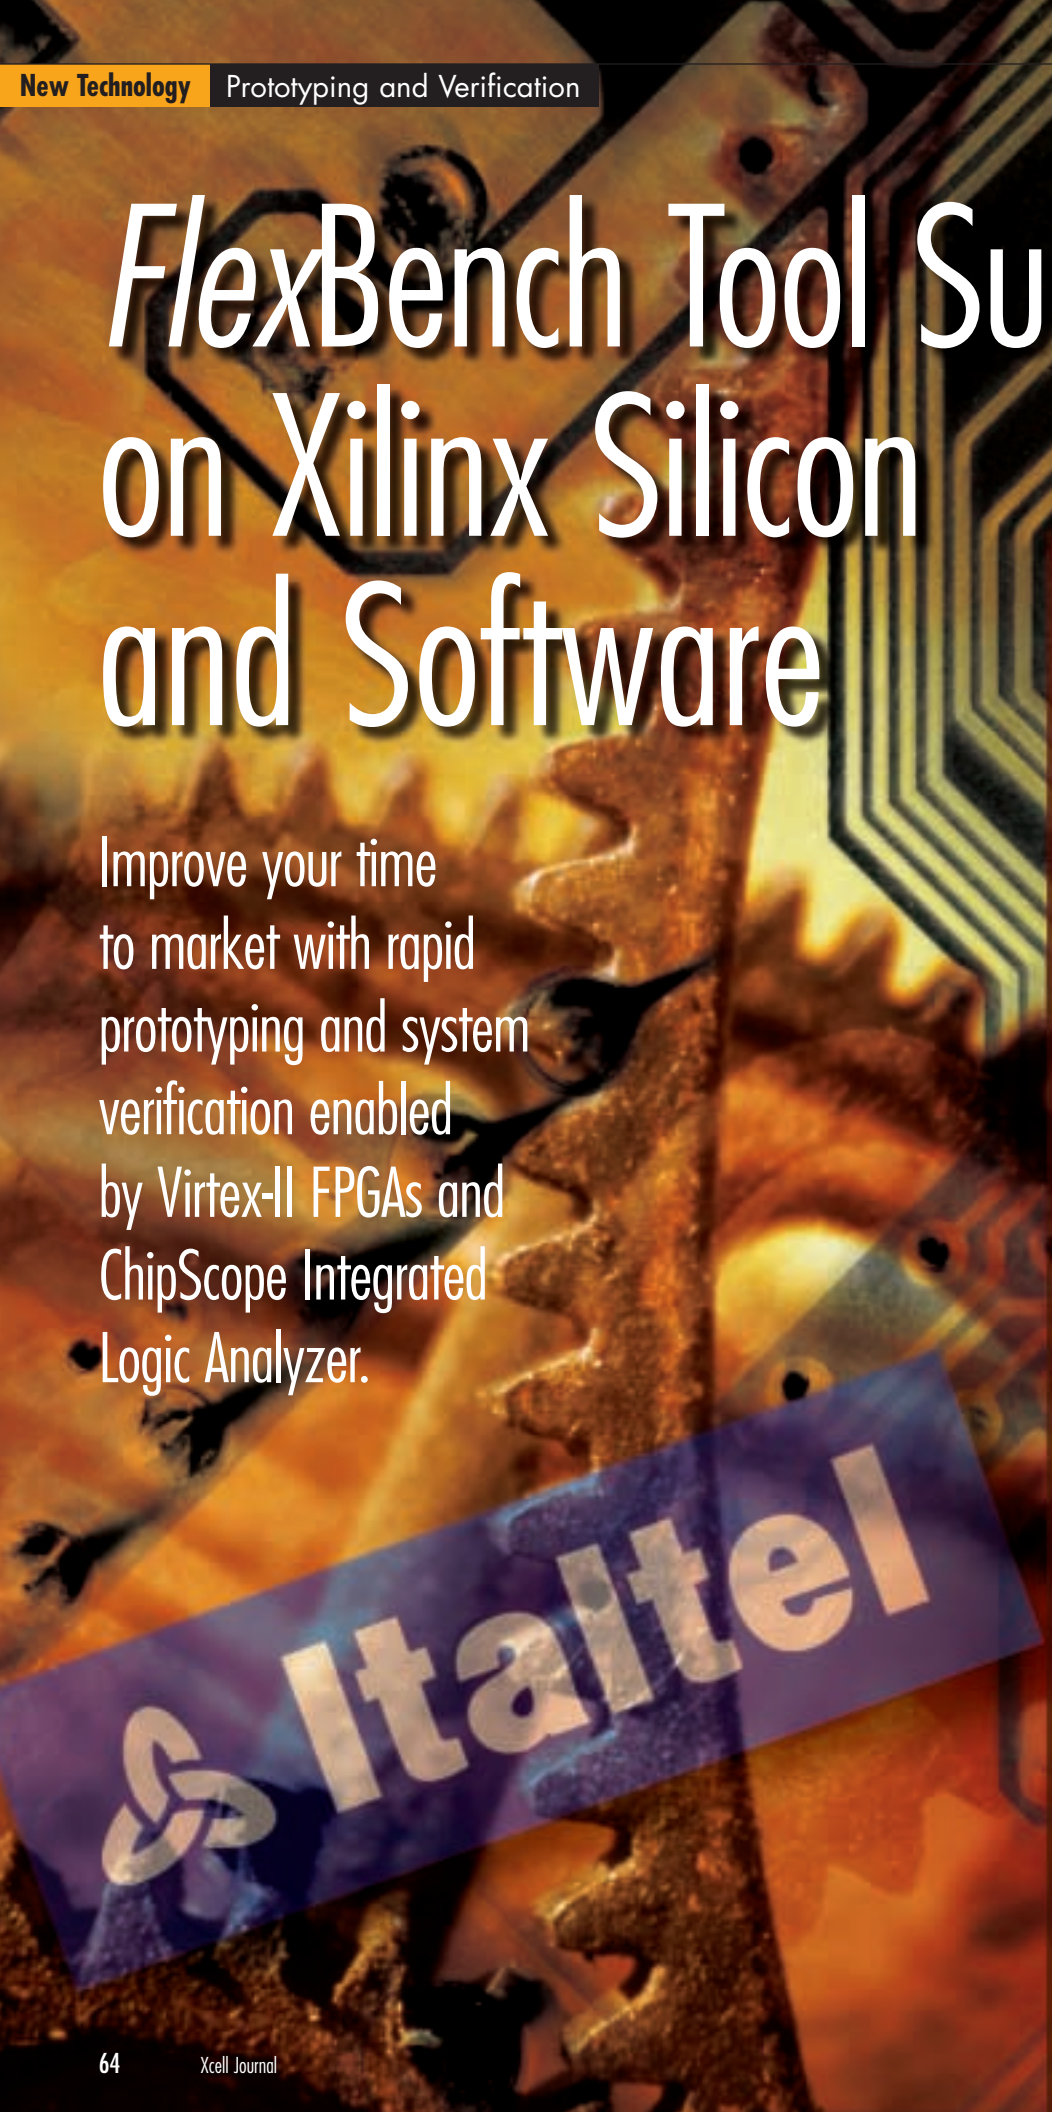

by Marco Pavesi
Head of Innovative Designs Laboratory
Italtel SpA
Marco.Pavesi@Italtel.it

System verification is the major bottleneck in the development of contemporary and future system-on-a-chip (SoC) designs. Increasing numbers of applications that process large quantities of data in real time (such as telecom and video) require verification techniques that run at or near real-time speeds. Delaying your software development until working devices are available is not a viable option in the face of enormous competitive time-to-market pressures.

Thus, early hardware/software (HW/SW) co-verification is not just practical – it's essential. At Italtel SpA, we and our partners have developed a superior method of improving system verification confidence through high-speed register transfer level (RTL) prototyping and the Virtex™ family of FPGAs.

Prototyping at Real Time Speeds

RTL prototyping allows you to run system hardware and software at speeds high enough to hunt and find hidden bugs. Just as important, you can confidently evaluate subjective characteristics, such as audio and video quality.

RTL prototyping uses off-the-shelf FPGAs to implement SoC custom logic, as well as test physical, real devices like memories, interfaces, and processors that comprise the other parts of the system to be verified. This strategy enables us to create, in effect, a clone of the SoC and all the parts of system that it interfaces with. We call this

assembling a "demonstrator." With a demonstrator, we can map the hardware and run the application software. RTL prototyping can be roughly divided in two types, custom and modular:

- Custom prototyping is the fastest technique, in terms of frequency. FPGAs and other devices are assembled on a board expressly designed for the demonstrator to be verified. Such a demonstrator may reach speeds as high as 200 MHz – but it requires several months from initial system design to the end of the verification process because of the intrinsic delay of the board fabrication process.
- Modular prototyping is not quite as fast as custom prototyping for verification – but when it comes to giving you the time-to-market advantage, it is, by far, the better solution. With modular prototyping, you can assemble FPGAs and other devices on general-purpose daughterboards that allow you to create a wide range of different demonstrators. Modular prototyping is a HW/SW co-development scheme based on a set of configurable carrier boards where daughterboards can be inserted and interconnected as necessary during the co-design process.

Avoiding the FPIC Dead End

Modular RTL prototyping platforms have intrinsic speed limitations related to modularity, accessibility, and routability. Field programmable interconnect chips (FPICs) have been the traditional solution for routability problems. Unfortunately, the existing technology trends for FPICs are insignificant (in terms of speed and the number of I/Os) compared to the skyrocketing speed and size of Virtex-II and Virtex-II PRO™ FPGAs. In short, no modular rapid prototyping platform based on FPIC technology can meet your customers' need for speed and time to market.

To approach real-time system speeds, a modular rapid prototyping tool with new interconnection technologies and a novel topology had to be developed. Such a tool – with single-ended, point-to-point signals – would have a physical speed limit ranging from 80 MHz to 100 MHz.

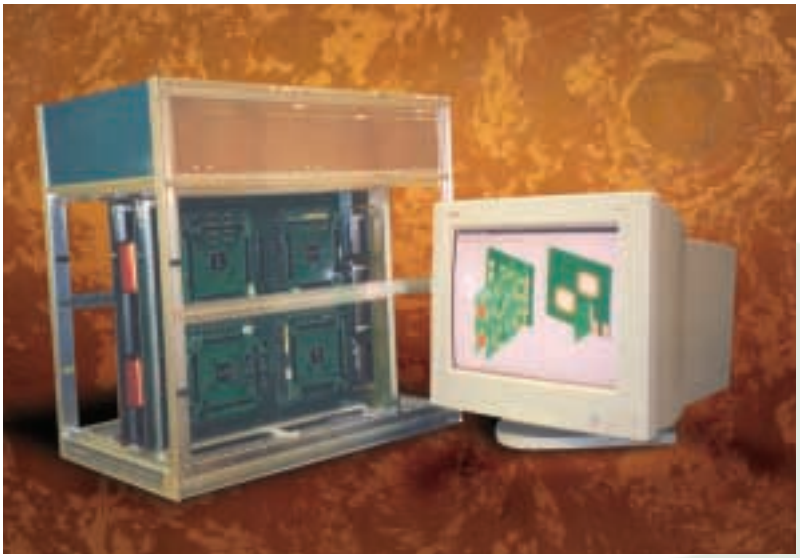

Figure 1 - Co-verification solution: Italtel FlexBench hardware with Temento Diaflex software

Given these parameters, three European companies collaborated for two years to design and develop the *Flex*Bench™ modular rapid prototyping tool suite:

- Italtel SpA, the largest Italian telecom manufacturer, is the leader for hardware development.
- Mistel SpA, another Italian telecom manufacturer, supports the hardware effort.
- Temento Systems, based in France, is dedicating its electronic test automation (ETA) software resources to the testing of SoCs and electronic boards.

Moreover, Oktet Ltd., a design service company based in St. Petersburg, Russia, was also deeply involved in the development the *Flex*Bench verification system.

Putting Together *Flex*Bench Hardware

As the *Flex*Bench project leader, I had been searching for a practical means of modular rapid prototyping in a HW/SW co-design environment. In 1999, two new technologies emerged that enabled our team to create and patent the *Flex*Bench concept:

1. QuickSwitch™ bus switches from Integrated Device Technology (IDT), Inc.
2. Mictor™ high-speed connectors from Tyco Electronics.

These two technological innovations gave our team the means of reconfiguring multiple printed circuit boards (PCBs) as needed during the HW/SW co-design process. Employing the reliable and fast-growing Xilinx FPGA technology, we envisaged it was possible to design the unique *Flex*Bench HW/SW tool suite.

The *Flex*Bench challenge was to create a rapid prototyping tool so general in purpose as to become the industry standard. Taking such a concept and turning it into a high-powered tool, however, is quite an undertaking that requires strategic alliances and adequate funding.

We applied for – and received – start-up funding from the European Commission. With that funding, we worked in a frenetic manner for two years to gather the expertise and to execute the rapid prototyping tool we envisioned.

At the hardware level, the *Flex*Bench rapid prototyping platform was to be basically a set of complex board designs. Fortunately, we had the PCB know-how and related resources in-house at Italtel – and valuable contributions from Mistel.

On the software side, the development of the computer assisted engineering (CAE) software we needed was not our area of competency. Therefore, we entered into a partnership with Temento Systems, a company known for its excellent ETA software. Temento developed the DiaFlex™ software tool we needed to monitor the *Flex*Bench system in action. See Figure 1.

Our acquisition of Certify™ RTL partitioning software from industry-leading Synplicity™ Inc. completed our primary software tool set.

Now that we had the boards and the software, we had to find the best programmable logic devices. It was a short search. In a class by themselves, Virtex-II Platform FPGAs from Xilinx perfectly fit the needs of our verification engineers: high speed, fast and simple compilation, and high capacity.

You can see the *Flex*Bench design chain in Figure 2.

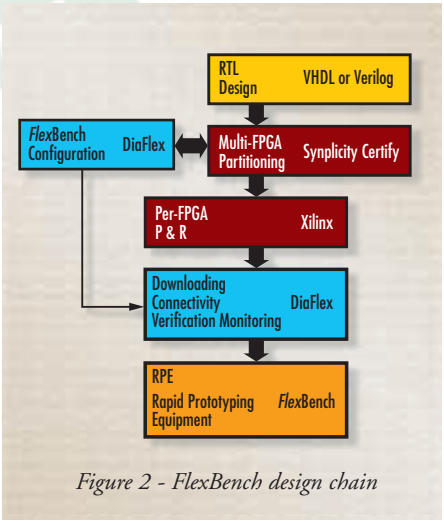

Figure 2 - FlexBench design chain

*Flex*Bench Rapid Prototyping at 100 MHz

As we said earlier, to achieve the speed edge in RTL prototyping with a modular tool, the FPIC approach is simply not suitable. A novel technology had to be introduced, based on the speed of IDT's QuickSwitch bus switches. By using pass-transistors, you can select interconnections among neighboring HW modules on multipurpose panels.

We built QuickSwitch bus switches and receptacles into the *Flex*Bench motherboards, which are populated by daughterboards named FlexPlugs. These modules are based on an open standard form factor. Using advanced, off-the-shelf interconnect technology, FlexPlug modules make available up to 888 functional signals. FlexPlugs host active devices, such as the leading edge Virtex-II 6000 Platform FPGA. The FlexPlugs also host power converters. Memories are allocated on small modules, named MiniPlugs, directly inserted into FlexPlugs in order to minimize interconnect delays.

We obtained the *Flex*Bench speed characteristics by interconnecting different FlexPlug modules through the distributed network of pass-transistors. Each side of a FlexPlug contacts the side of the nearest other FlexPlugs by means of channels. Channel size is configurable, through the JTAG port, from 0 to 222 functional wires with 8-bit granularity. Trace length minimization is achieved through a clever three-dimensional architecture.

The relatively short trace length of QuickSwitch delays (0.1 nanosecond) and the best-in-class set-up and clock-to-output parameters of Virtex-II FPGAs permits us to achieve 100 MHz clock speed (and over) on modular systems populated by up to 18 FPGAs.

The *Flex*Bench rapid prototyping equipment (RPE) is composed of a rack, a backplane (FlexPanel), a software-controlled clock generator (FlexClock), some carrier boards (FlexMothers), and several modules (FlexPlugs and MiniPlugs). FlexMother boards are connected through flexible printed circuits (FlexCable) and through the FlexPanel. See Figure 3 to see how the whole Flex hardware family interacts to form the *Flex*Bench verification system. Figure 4 shows a populated FlexMother.

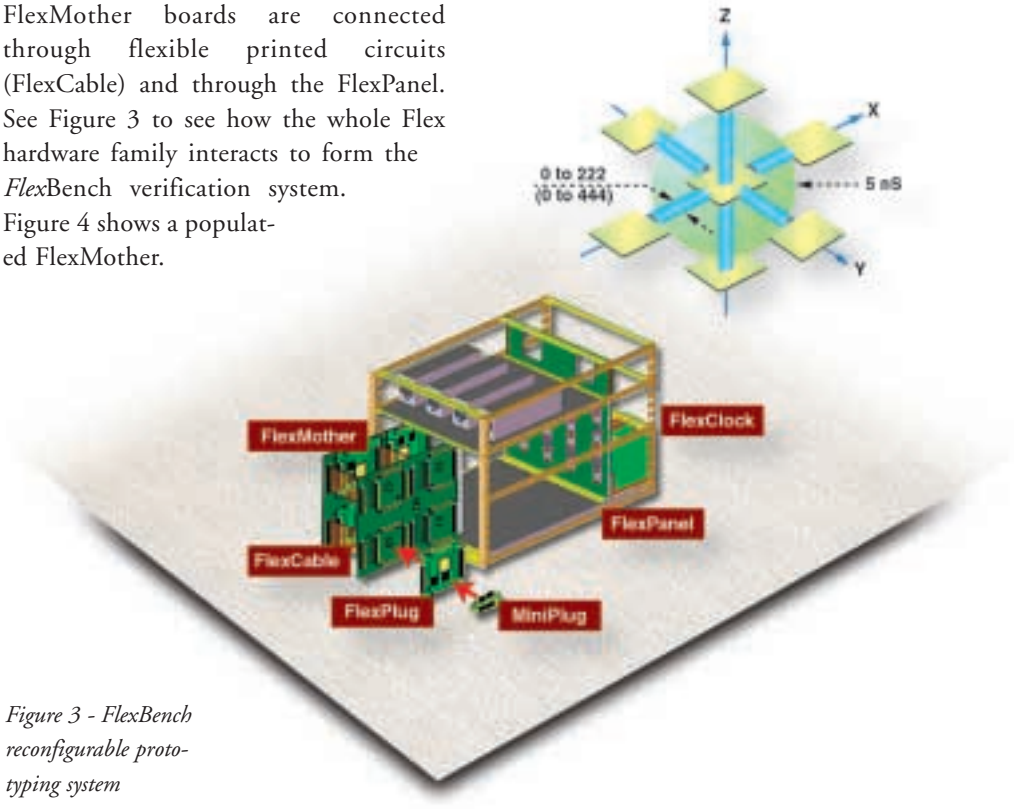

Figure 3 - FlexBench reconfigurable prototyping system

Integrating DiaFlex Software

Complementing the *Flex*Bench hardware is the DiaFlex software tool developed by Temento Systems. Dedicated to configuring, downloading, verifying, and controlling the *Flex*Bench RPE, the DiaFlex program suite allows you to design the RPE configuration through an intuitive, three-dimensional graphical user interface. Starting from RPE configuration, the DiaFlex application generates a flattened description of the platform definition into a Verilog format, called a VB

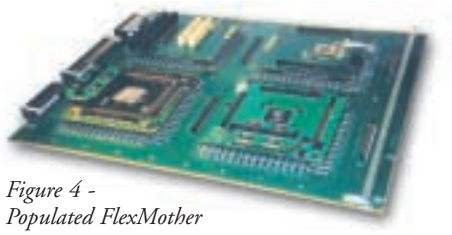

Figure 4 - Populated FlexMother

file. The VB file provides the Certify verification software with an easy view of the RPE resources and connectivity. Moreover, it provides direct channel configuration bitstreams .

The Certify program compiles RTL code, and spreads it according to the VB file description of the RPE. The Certify tool also

synthesizes all FPGAs. The FPGAs are then placed-and-routed via proprietary tools to produce FPGA configuration bitstreams.

The DiaFlex application downloads all bitstreams to the RPE and provides interactive debugging through a TemTag™ JTAG PCI board. Moreover, the DiaFlex software provides diagnostics, reports errors, and analyzes test results. This allows debugging at the signal name level, functional testing, and automatic management of IEEE 1149.1 test bench generation. Additionally, the DiaFlex program takes complete control of the FlexClock board. This enables you to configure FlexClock parameters: the frequency of clock synthesizers, the selection of global clock sources, and the control of other timing functions.

A Flexible Prototyping Library

To perform effective rapid prototyping, you must have a set of FlexPlugs and MiniPlugs. These plugs enable you to interconnect state-of-the-art FPGAs, memories, I/Os, processors, design platforms, and other components of your design.

While Italtel and our affiliates have designed general-purpose *Flex*Bench modules, design-specific modules must developed by final OEM customer. This is where the services of a company like Oktet Systems can become essential.

Sixteen general-purpose modules currently comprise the *Flex*Bench library. Some of the modules are displayed in Figure 5.

A Matter of Observability

Compared with the "observability" you can get with emulators, the classic argument against rapid prototyping has been that you can't "see" what's going on inside the design, because invasive physical probing is difficult, if not impossible, to accomplish in modular equipment.

Figure 5 - FlexBench library

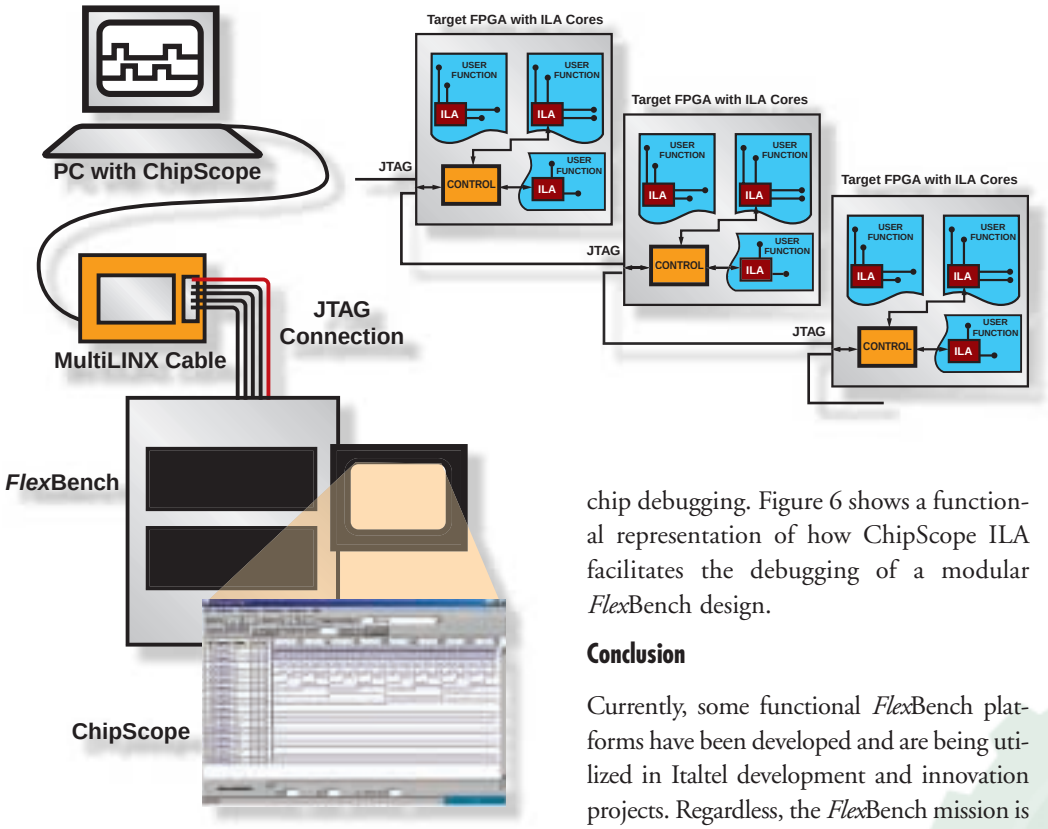

Figure 6 - ChipScope and FlexBench solution

Xilinx has removed this weakness thanks to its Integrated Logic Analysis (ILA) core, a solution that provides trigger and trace capture capability within the FPGA itself. Implemented by the Xilinx ChipScope™ Analyzer, the ILA core allows real-time access to any node in the *Flex*Bench chips. An easy-to-use graphical user interface provides the same functional controls of a sophisticated logic analyzer. With ChipScope ILA, you spend less time verifying chip functionality, and therefore, speed up your time to market.

By means of the ChipScope ILA, you get full visibility into the *Flex*Bench RPE. ChipScope ILA lets you see every selected node in every FPGA used in your design. This virtually eliminates the need for invasive physical probing of the *Flex*Bench RPE. With the Xilinx ChipScope ILA solution, you can perform real-time, on-chip debugging. Figure 6 shows a functional representation of how ChipScope ILA facilitates the debugging of a modular *Flex*Bench design.

Conclusion

Currently, some functional *Flex*Bench platforms have been developed and are being utilized in Italtel development and innovation projects. Regardless, the *Flex*Bench mission is not to become a proprietary tool, but to become an industrial standard.

Italtel/Temento are and will be engaged in trials with major silicon vendors to demonstrate and to prove the *Flex*Bench technology. We believe the versatile *Flex*Bench model is a major step in rapid prototyping and system verification. We are proud to have designed it. To ensure best diffusion of this tool, we are now researching the best sales channel.

Multiprotocol Modular Engineering Solutions Platform Offers Design Flexibility and Faster Time to Market

The Nu Horizons Engineering Solutions Platform presents a cost-effective alternative to single-board tools for developing high-end data communication and server applications.

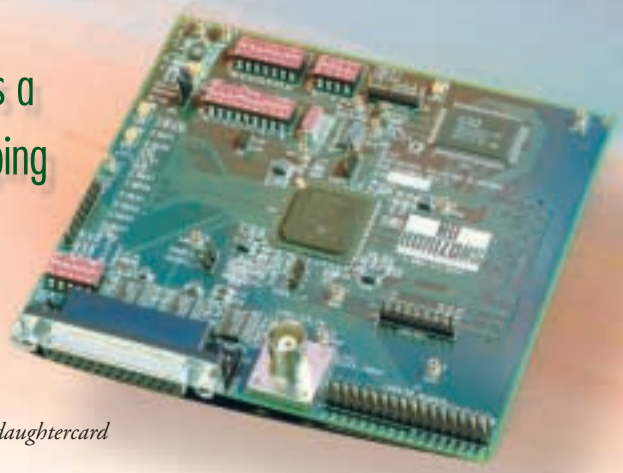

Virtex-II daughtercard

by Bill Pratt
Regional Applications Manager
Nu Horizons Electronics
Bpratt@NuHorizons.com

Today's one-sided trend toward shorter system development cycles and increasingly complex designs are driving the need for more timely system validation. Although designers recognize this weighted ratio, and are turning more and more to Xilinx FPGAs for their designs, they have also recognized a need for development boards that are up to the challenge.

Nu Horizons Electronics Corp. understands this short-cycle/high-complexity design ratio and has introduced a line of hardware development platform tools that facilitate system verification within an integrated environment. The Nu Horizons Engineering Solutions Platforms are based on a modular design concept that allows the re-use of expansion cards with other interface expansion cards to create a unique development evaluation environment.

Modular Design Solution

The Engineering Solutions Platform solution consists of up to three boards in a modular environment. With this modularity, you may combine a number of different expansion cards for multiple solutions, or you can build a total solution in a modular environment. The days of purchasing a development board for a single application are gone.

Engineering Support Platform motherboard

Motherboard

The primary function of the Engineering Solutions Platform motherboard is to provide a high-speed backplane for board-to-board data transfers. This backplane is called the Horizon Bus. With the high-density board-to-board connectors, you can evaluate the functions of two expansion cards, which allows emulation of a total solution. The Horizon Bus has an aggregate bandwidth of 26 Gbps, and is highly suited for developing solutions for the telecommunications and data network markets.

The main board, or motherboard (Figure 1), has the following features:

- Xilinx XCR3128XL -10PQ VQ100 CoolRunner™ CPLD
- High-density board-to-board connectors for two expansion cards
- Dual seven-segment LCD display
- RS-232 level shifter for UART core instantiation
- JTAG configuration headers
- Two 44-pin PLCC sockets for Xilinx configuration PROMs
- Fifty test points
- Power management
- Reset controller
- Status LEDs.

Virtex-II Expansion Card

The Xilinx Virtex™-II expansion card (Figure 2) is a cost-effective development board with a broad suite of features. Available now is the XC2V1000 FG456 one-million-gate Virtex-II device with the following features as implemented on the Engineering Solutions Platform daughtercard:

- 16 channels of LVDS interface
  - – User-configurable as eight transmit and eight receive or 16 transmit channels
- Multiprotocol serial clock/data transceivers with auto cable termination
  - – Certified TBR-1, TBR-2, NET-1, and NET-2 compliant.
- 16 megabits of synchronous NBT SRAM
  - – Flow through and pipelined
- Programmable clock from 1 MHz to 200 MHz
- Six user clock inputs
- High speed 133-bit at 200 MHz backplane
- DCI (digitally controlled impedance).

Design Platform for High-End Applications

Targeted at high-end data communications applications, the Virtex-II daughtercard is applicable to many scenarios, including SPI-3 to CSIX bus interface or offering transparency within high-end networking solutions. The 16 channels of high-performance LVDS (low voltage differential signaling) provide the capability to prototype high-speed differential interfaces such as the SPI-3 and SPI-4 bus technologies. The Receive port includes parallel termination; the Transmit side does not require this feature.

The multiprotocol clock/data transceivers with auto cable termination can support multiple protocols such as two channels of V.35 HDLC (high-level data link control). Coupled with the HDLC core from Xilinx, you can turn your Virtex-II FPGA into a powerful network engine. The multiprotocol clock/data transceivers also support X.21, V.11, RS-232, RS-449, and RS-53W\ serial interfaces.

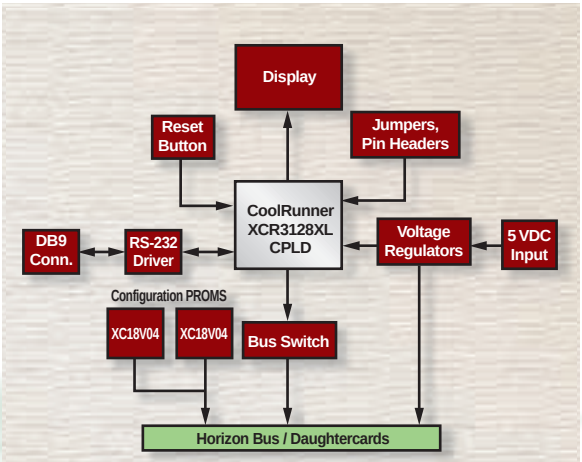

Figure 1 - ESP motherboard block diagram

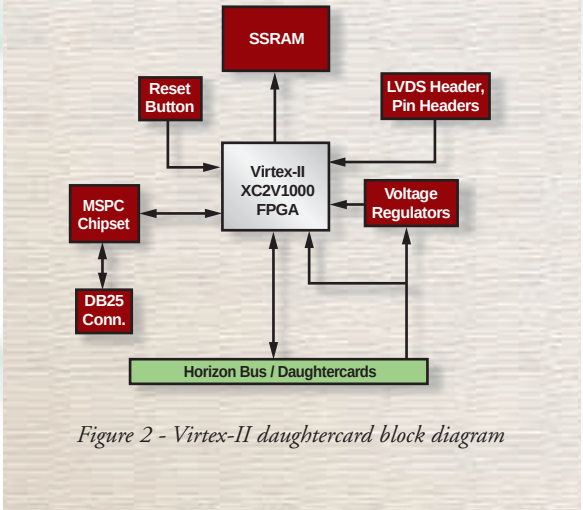

Figure 2 - Virtex-II daughtercard block diagram

NBT (no bus turnaround) RAM offers a standard packaging for high-speed SSRAM devices. The Virtex-II daughtercard has 16 MB of NBT SSRAM from GSI Technology, which allows for zero-wait, read-write bus utilization. This GSI SSRAM is programmable as FT (flow through) or pipelined, operating as a pipelined synchronous device. This means that in addition to the rising-edge triggered registers that capture input signals, the device incorporates a rising-edge triggered output register.

The daughtercard also incorporates an ICS525-02 user-configurable clock chip from Integrated Circuit Systems Inc. The ICS525-02 is the most flexible way to generate a wide range of highly accurate clock output from a standard crystal or clock oscillator. You can easily program the ICS525-02 to output a frequency from 1 MHz to 200 MHz by setting a bank of switches located on the board. You may also supply a LVDS clock, as well as four other single-ended user supplied clocks, for a total of six usable clock inputs.

With the Horizon Bus backplane, the Virtex-II daughtercard becomes a very powerful development platform offering up to 26 Gbps throughput and allowing the instantiation of a multitude of bus architectures, such as CSIX for telecommunications.

Nu Horizons Electronics Corp. is in the process of developing new daughtercards for the Engineering Solutions Platform. Future cards include an HTML server/white appliance controller and a high-speed AFE (analog front end) with an Ethernet interface.

Conclusion

The Engineering Solutions Platform from Nu Horizons Electronics Corp. is a modular Virtex-II development environment that provides designers with the hardware necessary to develop, prototype, verify, and test their designs before they are finalized. Intended to accelerate the design cycle timeline, the Engineering Solutions Platform is a modular solution that enhances and simplifies system design and validation processes. Its modularity makes it highly functional over many design cycles, thus providing a cost-effective alternative to single-board tools for developing high-end data communication and server applications. The Engineering Solutions Platform motherboard with the Virtex-II daughtercard sells for $999 and is available for immediate shipment. For more information, please log onto *www.nuhorizons.com.*

Amplify

Plug in your high-density FPGA.

And get ready. You've never seen performance like this. Introducing Amplify® Physical Optimizer™ — the first and only physical synthesis solution for programmable logic. Now you can achieve aggressive performance goals, and save weeks while you're at it. Easy to learn and use, Amplify delivers better Quality of Results by utilizing both physical and timing constraints during synthesis. As a result, designers are achieving up to 35% performance gains, on or ahead of schedule. Plus, Amplify supports and enhances Team Design. Not only does it manage the physical hierarchy of a design, Amplify optimizes performance, regardless of how a design is allocated across the team. Get in on the outstanding performance. For more information and an evaluation copy, visit **www.synplicity.com/amplify**.

Join Synplicity at Programmable World 2002, and attend our upcoming webcast seminar featuring Synplicity's Amplify Physical Optimizer and Xilinx® Virtex-II Pro™. Details at www.synplicity.com/events/xilinx_events.html

See us at the Xilinx Programmable World 2002

Synplicity®

Simply Better Results

www.synplicity.com

phone: **(408) 215-6000**
email: **info@synplicity.com**

Parallel Cable IV Connects Faster and Better

This new high-speed download cable supports ultra-low voltages.

by Theresa Vu
Product Marketing Engineer
Xilinx, Inc.
theresa.vu@xilinx.com

The new Parallel Cable IV downloads data more than 10X faster than the previous JTAG/Parallel Cable III. Now, you can download to one XC2V8000 Virtex™-II FPGA in less than eight seconds.

The PCIV features ultra low-voltage support for all Xilinx FPGA, CPLD, System ACE™ MPM (multi-package module), and ISP (in-system programmable) PROM devices (see Table 1).

The PCIV connects to any desktop or laptop computer using the standard IEEE 1284-compliant parallel port and draws power directly from the computer (through the mouse/keyboard port) or an external power supply.

A robust ribbon cable ships with the PCIV, which gives you the flexibility to use either the JTAG (IEEE 1149.1) or Slave Serial download mode at the fastest speeds. The small profile of the ribbon cable connector minimizes the need for board space. The ribbon cable offers an error-free, quick connect target interface compared to cumbersome flying lead wires of Parallel Cable III.

The PCIV is backward compatible with the PC III and offers a connector for flying lead wires.

PCIV extends Xilinx leadership in pre-engineered configuration solutions by offering a fast, simple, low-cost download solution for all Xilinx FGPA, CPLD, System ACE MPM, and ISP PROM devices. The Parallel Cable IV will be available in late March through Xilinx distributors and the Xilinx e-commerce site for $95. For more information, see *www.xilinx.com/support/programr/cables.htm.*

	Parallel Cable IV	Parallel Cable III
Download Speed	up to 4 Mbps	up to 300 Kbps
I/O Voltage Support	1.5V, 1.8V, 2.5V, 3.3V, 5V	2.5V, 3.3V, 5V
Download Modes	JTAG (IEEE 1149.1) and Slave Serial	JTAG (IEEE 1149.1) and Slave Serial
Power Supply	PC or External Power Supply	Target System
Board Connections	Ribbon Cable and Flying Wires	Flying Wires
Software Support	iMPACT	iMPACT

Table 1 - Parallel Cable IV vs. Parallel Cable III

Parallel Cable IV connected to PC

Parallel Cable IV connected to external power supply

Parallel Cable IV with ribbon cable connector

scc-*II* Microsequencer – A New Solution for Platform FPGA Designs

When your project design is too big for a finite state machine, but a microcontroller would be overkill, try Ponderosa Design's scc-*II* microsequencer.

by Aki Niimura
Consultant
Ponderosa Design
ponderosa_design@pacbell.net

As the complexity of FPGA-based systems grows every year, we are asked to implement larger, more complex functionality within tighter schedules. Furthermore, the type of design has changed rapidly in recent years. People used to design an FPGA taking an existing board design, often containing asynchronous clocks. Those days are over. You can not design today's Platform FPGA just by extending yesterday's design practices. New designs often require the implementation of complex sequences or communication protocols. The finite state machine (FSM) is a well-known design methodology to implement such sequences. FSM is very effective when the sequence is not very complex. However, implementing a complex sequence using an FSM is not practical and is often difficult to maintain.

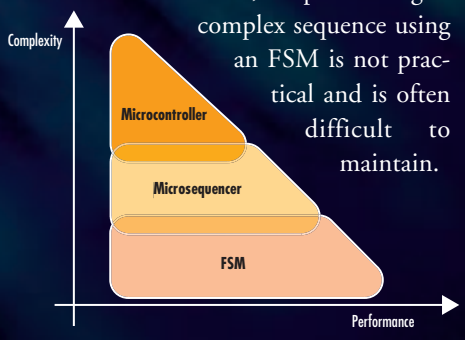

Figure 1 - Performance versus functional complexity

Microcontrollers are commonly used to implement complex protocols. However, they require substantial resources (memory, cost, pins, ...), which can be difficult to justify in real-life situations. On the other hand, software implementations allow designers to cope with mounting logic complexity. They are easy and quick to implement and easier to maintain.

There is a gap, however, between the range of design complexity that FSM methodology can handle and what microcontroller-based methodology is good for, as shown in Figure 1. As microcontrollers become more powerful, the gap is widening.

The scc-*II* is not just another set of microcontrollers. We specifically constructed the microsequencer to fill the gap between low-level FSM solutions and high-level microcontroller designs.

scc-*II* – A Configurable Microsequencer

Sequencers have been used in many LSI projects to implement functions. For example, instructions in a CISC (Complex Instruction Set Computer) CPU were often implemented in this way (called microcode, which is written in a proprietary assembly language). By allowing users to write programs in a high-level language, the scc-*II* can accommodate a wider range of FPGA applications.

The key architectural benefits of the scc-*II* are:

- Small footprint
- High-level language support
- Small code size
- Configurable and customizable
- Capable of handling 16-bit and 32-bit data types
- Timer (integrated into the core architecture)
- Support of interrupt handling
- Developing and debugging tools
- Utilization of Xilinx Spartan™-II and Virtex™-II devices.

A block diagram of the scc-*II* is shown in Figure 2. The core itself requires 400 to 600 LUTs, depending on the configuration and synthesis constraints.

How the scc-*II* Works

The scc-*II* employs a stack-based architecture. Stack computers use data stacks to evaluate given operations (Figure 3). The benefits of stack-based architecture are:

- High-level language ready – can execute syntax tree directly
- Simple hardware – easy to understand, easy to customize
- Small instruction code – most scc-*II* instructions are one byte long.

Another unique aspect of the scc-*II* architecture is the use of register windows. Register windows are used to pass arguments to a function being called. Because the scc-*II* does not use a stack frame in memory to pass arguments, the scc-*II* does not require data memory to run a high-level language program, thus making the scc-*II* more attractive for Platform FPGA applications.

Programs for the scc-*II* are almost entirely written in the high-level language SC.

The Language SC

The scc-*II* assumes the use of a high-level language. However, existing high-level languages are not designed for microsequencer applications. Therefore, we developed a stripped-down version of C language – SC. SC programs do not support "struct" and other complex data types, but SC has several enhancements to describe control applications efficiently. Timer, I/O, and debug features are natively supported in SC.

The following is a code fragment from a project that controls an SDRAM memory.

```
void
init_sdram()
{
   while(!eval_cond(DLL_RDY)) { }    // wait for DLL is ready

   wait(14);      // 286 uS
   outp(SDCMD, SD_PRE);
   repeat(8) {
    outp(SDCMD, SD_AREF);          // looping takes 7 cycles
   }          // tRC = 84nS; 20nS * 5
   outp(SDCMD, SD_MODE);
}
```

In the above code fragment, eval_cond(n), wait(n), and outp(port) are not function calls, but they are natively supported by SC. Note that the loop counter of the repeat statement is placed in the data stack and not in the register file.

scc-*II* Target Applications

The scc-*II* can be used in designing functional blocks to perform procedural control. For example, flow charts or simple arithmetic functions (such as averaging) are good candidates for implementation by the scc-*II*.

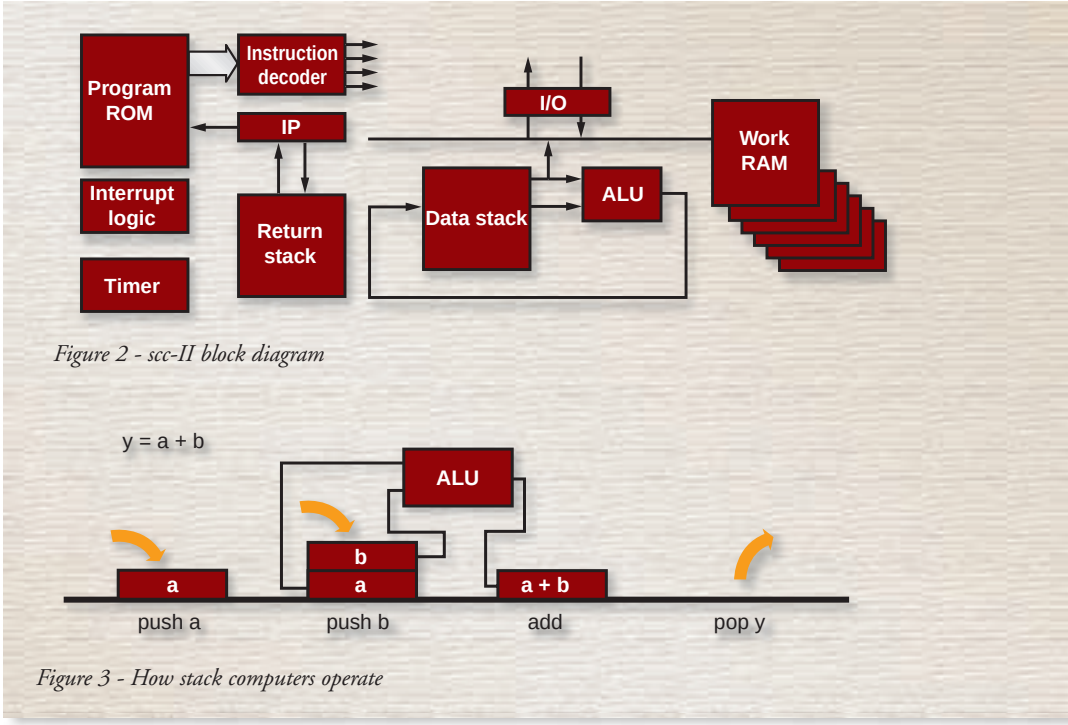

Figure 2 - scc-II block diagram

Figure 3 - How stack computers operate

Other target applications include:

- Read/write flash, EEPROM, 1-Wire™ devices
- Interface to I2C, RS-232, Ethernet, USB1.1, IrDA
- Command interpreter (a block is controlled through commands)
- User interface (such as keypad, LCD, touch panel)
- Servo controller (some arithmetic operations required)
- Design that requires many variants.

The scc-*II* and Virtex FPGAs

The Virtex family of FPGAs are true "system on a chip" platforms. The advanced technology available in Virtex-II devices provides further attractive features to the scc-*II*, including:

- The scc-*II* can run at 70 MHz or faster.
- Larger Block RAM allows larger program sizes (up to 8 KW).
- Native multiply operators are supported.

Case Study – Web on FPGA

To demonstrate the effectiveness of the scc-*II* solution, we have developed a Web server that uses less than 25% of the resources of a Spartan-II FPGA (XC2S150). The only additional hardware required beside the FPGA are an Ethernet PHY device and a signature ROM (optional). We found that the Spartan-II VoIP Development Kit from Insight Electronics *(www.insight-electronics.com)* included all the hardware components we needed. Thus, we decided to use this off-the-shelf board to create our Web server design. Figure 4 shows a screenshot from a Web server proj-

Figure 4 - WebThermo demo application

ect called WebThermo. The screen displays the current temperature every minute. Table 1 shows utilization statistics from a Synplify analysis of the XC2S150 device.

WebThermo	Usage of XC2S150	Note
LUTs	828 (23%)	Synplify 7.0
Block RAM	8 of 12	3 for TX, RX buffer, 5 for Program
BUFTs	320 (18%)	

Table 1 - WebThermo logic size

The 2 KB WebThermo program implements all Ethernet, TCP/IP, and Web (HTTP) protocols, as well as Celsius-to-Fahrenheit conversion. At start up, the program retrieves a unique 48-bit ID code from a Dallas 1-Wire device (DS18S20), which is used as an Ethernet MAC address. For further details on the WebThermo project, please visit home.pacbell.net/akineko/.

Program Development

One challenge of the scc-*II* solution is in providing reasonable program development and debugging tools. Figure 5 illustrates a typical program development flow. In addition to key software tools, we wrote many scripts and templates to automate the design process. While creating several projects with the scc-*II*, we refined the RTL design, as well as the development software and scripts. As a result, they have become mature and stable.

Currently the development environment is supported under Unix. It is also possible to port some of the tools to Windows platform using Cygwin from Cygnus (RedHat). The tools are developed assuming that the user's RTL design is in Verilog HDL.

Debugging, Then Debugging Again

Debugging is the biggest challenge in developing an scc-*II* based design. We are providing several debugging aids:

1. Debugging starts with simulation:

– Three debug instructions (print, $dump, $stop)

– Self-checking embedded in the code

– Execution trace log generation

– Dis-assembler to display current context (on-the-fly/offline)

2. Ready to try on the board:

– JTAG debugger to download program without backend (synthesis + PAR)

– UART customized for debugging (one can use printf())

– "xdl" script to replace ROM contents without backend.

Lessons learned:

1. Logic simulation is always the best tool for debugging.
2. printf() is a primitive but very powerful means for debugging.
3. Use #ifdef ... #else ... #endif to switch between debug and release.
4. A bigger vehicle is needed for debugging (you may need 2 KB to develop a 1 KB program).

JTAG debugger

The JTAG debugger (jtagdbg) has proved to be a powerful tool to facilitate the debugging process. The JTAG debugger uses the Virtex USER1 JTAG command to communicate with a Virtex FPGA. By substituting the instruction ROM block in the scc-*II* design with a JTAG embedded ROM block, you can perform several debug commands, such as downloading a program without going through the FPGA backend design process. No signal change is required, as JTAG signals are hidden from your RTL code.

Conclusion

We have presented a microsequencer, the scc-*II*, which is new to conventional FPGA design practices. Unlike other IP cores, the potential of the scc-*II* is not limited to its original form. Rather, the scc-*II* can evolve to meet each application challenge. One avenue we plan to explore is adding Galois instructions to the original scc-*II* core. This enhancement can help in error correction or security applications.

Another avenue we plan to pursue is project automation, such as a wizard script that sets up project directories and tools – and then creates a skeleton version of RTL code, as well as skeleton SC program and header files.

The complete scc-*II* design solution is offered by Ponderosa Design in Sunnyvale, California Please write *ponderosa_design@pacbell.net* or visit *http://home.pacbell.net/akineko/.*

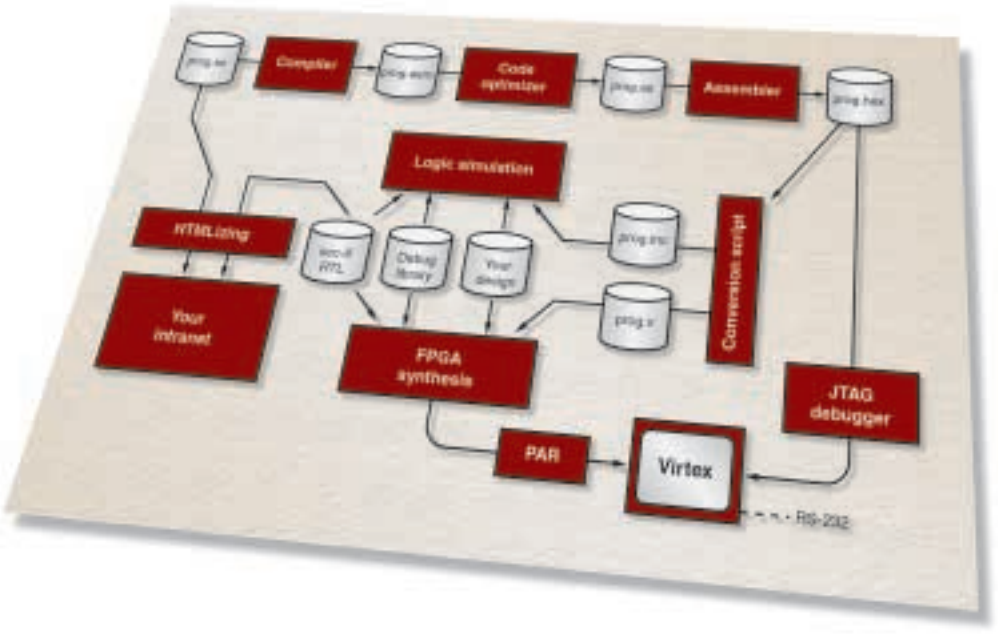

Figure 5 - The scc-II program development flow

Upgrade to Synopsys FPGA Compiler II Synthesis Tool to Maximize Virtex-II Pro Performance

FPGA Compiler II's unique algorithms aid in designing chips correctly and on time.

by Jackie Patterson
Director of Marketing Programs
Synopsys, Inc.
jackiep@synopsys.com

When you target Virtex-II Pro™ Platform FPGA, you are using one of the best FPGAs on the market – and you need a high-performance synthesis tool to match. That's why Synopsys Inc. is focusing efforts on its premier FPGA Compiler II™ FPGA synthesis tool and continuing to hone its support for top-of-the-line programmable logic devices. In 2002, FPGA Compiler II is superceding FPGA Express™, which will be discontinued.

This article will explain the differences between the two synthesis tools, summarize key features carried forward from FPGA Express to the FPGA Compiler II, and explain how you can upgrade to FPGA Compiler II today.

FPGA Compiler II Has Unique Capabilities

FPGA Compiler II (FCII) was developed for high-performance devices such as Virtex-II Pro Platform FPGAs. FCII combines the architecture-specific synthesis engine of FPGA Express with advanced technologies suitable for ASIC-like design challenges, as shown in Figure 1. These leading-edge capabilities are unique to FPGA Compiler II and not found in any other synthesis tools on the market.

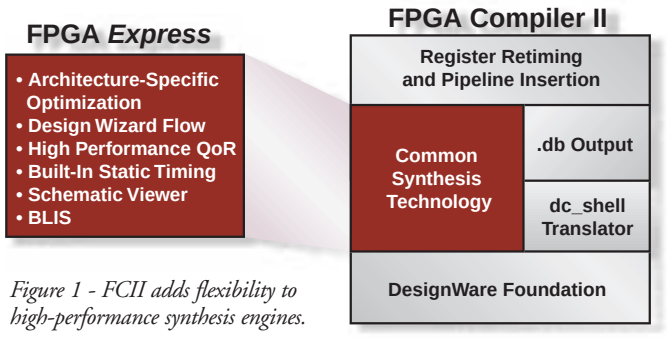

Figure 1 - FCII adds flexibility to high-performance synthesis engines.

Register Retiming

FCII uses sophisticated retiming algorithms to boost clock speed automatically. Retiming works by analyzing all the combinational logic in the design, and then selecting the optimal register placement to meet your design goals. See Figure 2.

Retiming can also pipeline your design automatically. Just code in the number of register banks to match your latency requirements, and FCII does the rest, moving the registers into the optimal position to form your pipeline.

Register retiming is easy to use in FCII. All you have to do is set the retiming variable in your script, and FCII will automatically position your registers in the optimal places to maximize clock speed.

```
current_chip my_chip
set_chip_retiming -enable
```

You can also use the FCII graphical interface to invoke retiming, if you prefer. Just select Edit->Constraints from the menu and check the retiming box in the Xilinx vendor options tab. Synopsys full-chip retiming optimizes performance of all datapath, control, and random logic in the high end of the Xilinx line – the Virtex™, Virtex-II, and Virtex-II Pro families of devices.

DesignWare Foundation IP Library

ASIC designers save design time by reusing components from the Synopsys DesignWare Foundation IP library. Although some components in the library are also useful for production FPGA designs, the biggest benefits come for those who are using a Virtex-II Pro device

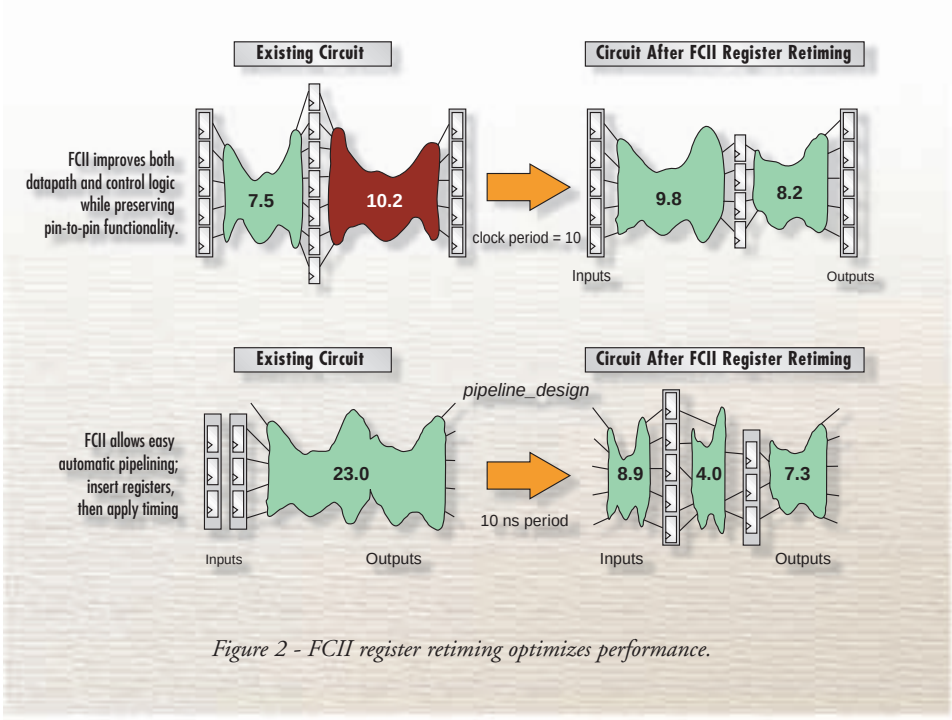

Figure 2 - FCII register retiming optimizes performance.

to prototype an ASIC or SOC. For a prototype, FCII's ability to implement the DesignWare components in the FPGA provides confidence that the ASIC will work as planned, because it allows you to verify exactly the same IP that is used in the ASIC.

Full Flow Support

FPGA Compiler II allows designers to take advantage of other high-performance tools for FPGA design and verification. When undertaking large, complex Virtex-II Pro devices, designers need to verify the silicon before it gets into the lab. That process starts with a quick analysis of the source code using the LEDA® HDL checker. The LEDA checker leverages general-purpose rules along with Xilinx-specific rules to alert designers to potential problems and optimization opportunities in the HDL. Then the HDL functionality is verified using a fast simulator such as VCS™ for Verilog or Scirocco™ for VHDL. After the optimization by FPGA Compiler II, designers can use Formality® to formally verify the design, avoiding time-consuming simulation runs. Of course, the timing of the design is comprehensively checked through static timing analysis. In this way, Synopsys FCII customers can leverage all available means to get their designs done correctly and on time.

Architecture-Specific Synthesis Engine

FPGA Compiler II carries forward the same architecture-specific synthesis engine you've relied on in FPGA Express. This includes:

- Support for all Xilinx devices through Virtex-II Pro and beyond
- Block Level Incremental Synthesis (BLIS) to speed your design cycle time by re-doing only the blocks in the design that have changed
- User control of register duplication to eliminate critical paths caused by high fan-out nets
- Networked licenses on both UNIX and PC
- Integrated schematic viewing and timing analysis
- Full TCL scripting
- ROM inference.

Upgrading to FPGA Compiler II

FPGA Compiler II is the choice for synthesis of high-performance FPGAs such as the Virtex-II Pro pf1.

To upgrade to FPGA Compiler II, contact your Synopsys account manager or visit *www.synopsys.com/fpga*.

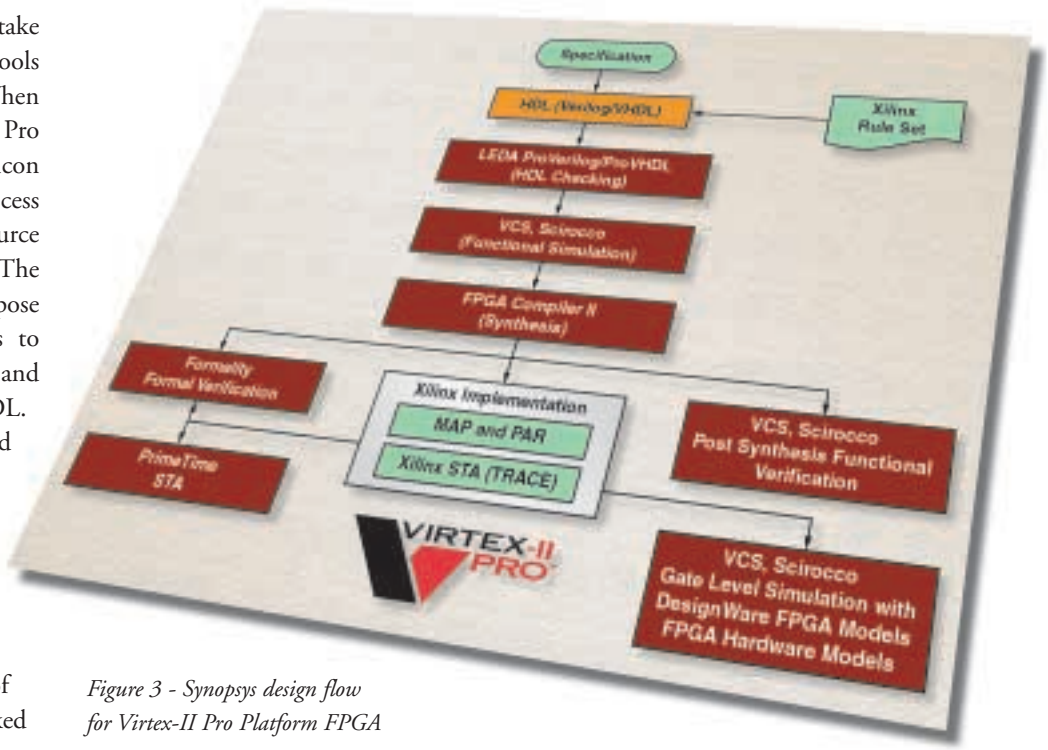

Figure 3 - Synopsys design flow for Virtex-II Pro Platform FPGA

Use Rocket I/O Multi-Gigabit Transceivers to Double Your FPGA Bandwidth

Virtex-II Pro Platform FPGAs break open the I/O bottleneck.

by Brian Von Herzen, Ph.D.
CEO, Rapid Prototypes, Inc.
a member of Xilinx XPERTS Program –
Xilinx Worldwide 3rd Party Certified Design Centers
Brian@FPGA.com

As FPGAs increase in size and performance, I/O resources become the main bottleneck to FPGA performance. Although the effective area of a chip grows as the square of the feature size, the perimeter I/Os grow only linearly. State of the art designs require higher performance I/O modules.

In response to this increasing demand on I/O resources, Xilinx has developed novel I/O structures called Rocket I/O™ multi-gigabit transceivers (MGTs) that enable order-of-magnitude increases in I/O performance. The Rocket I/O MGTs double the total I/O bandwidth of the Virtex-II Pro™ family of devices using only a few percent of the pins.

With up to 16 MGTs per device, the Virtex-II Pro achieves an additional 100 gigabits per second of I/O bandwidth in the larger devices over what is available with the general-purpose I/O blocks. Rocket I/O MGTs enable multiple gigabit I/O standards and maximize performance for FPGA-to-FPGA communications. Even though Rocket I/O MGTs dramatically increase performance for demanding applications, they are easy enough to use for simple FPGA-to-FPGA communications with special soft macros such as the Aurora core available from Xilinx. The interface has been simplified to the extent that no external resistive termination is required with the Rocket I/O MGTs. The transceivers can be internally configured to match 50Ω or 75Ω transmission lines.

MGTs Onboard

The Rocket I/O MGTs are shown in Figure 1, which illustrates the overall Virtex-II Pro architecture. The MGTs are located above and below columns of Block RAM, providing close availability of Block RAMs for ingress and egress FIFOs. As many as 16 MGTs are integrated on each FPGA above and below the Block RAM columns. The clock distribution networks can feed these transceivers for low-skew clock alignment between MGTs.

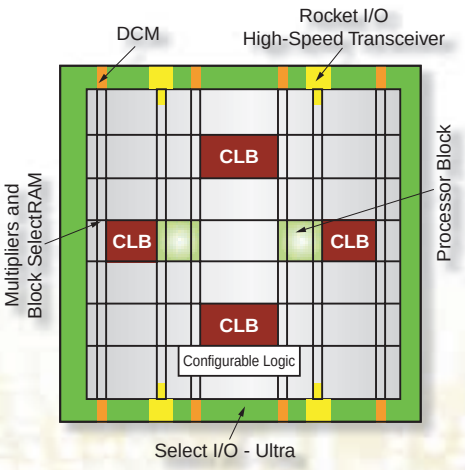

Figure 1 - Virtex-II Pro FPGA architecture, including up to 16 Rocket I/O MGTs

MGTs = Multiple Gigabit Standards

The Rocket I/O MGTs have been designed to be compliant with:

- Gigabit Ethernet
- 10 Gigabit Ethernet XAUI
- Fibre Channel
- InfiniBand™ Architecture
- Xilinx Aurora core

Configurable hardware support is provided for:

- 8B/10B encoding
- Disparity control
- Transmitter and receiver termination impedance
- Pre-emphasis
- Amplitude control
- Loopback testing.

These are the essential configurable hardware features that enable compliance with all of the main multi-gigabit signaling standards. You can find more information on Virtex-II Pro standards support on the Xilinx website at www.xilinx.com/partinfo/databook.htm.

MGTs Speed up FPGA Communications

Although MGTs have many important applications interfacing to industry standard gigabit communications protocols, they can serve another important function in boosting the bandwidth available among FPGAs.

Figure 2 shows a data communications application requiring two FPGAs. For a typical implementation, half of the I/O bandwidth is allocated to external links and half to the inter-FPGA link. With 800 general-purpose I/O pins available, running at a typical speed of 250 MHz, the pair of FPGAs can support a pipeline throughput of 100 gigabits per second.

If the 16 bidirectional MGT links are added, 16 links x 3.125 gigabits are available in each direction, for another 100 gigabits per second total. Figure 3 shows

Figure 2 - FPGA processing pipeline using general purpose I/O (GPIO) only

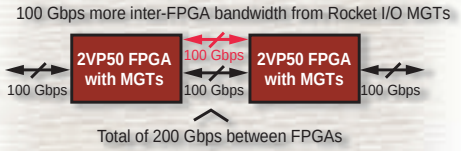

Figure 3 - FPGA processing pipeline using Rocket I/O multi-gigabit transceivers and general purpose I/O (GPIO)

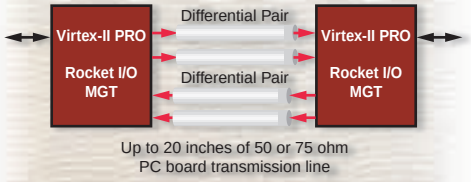

Figure 4 - Rocket I/O MGTs configured to pre-emphasize data transmission over PCB traces up to 20 inches long

the system diagram for two Virtex II Pro devices including the Rocket I/O MGTs between the devices. With Rocket I/O, the total bandwidth available between the two FPGAs increases to 200 gigabits per second. This extra bandwidth is extremely useful in switching applications that may require up to 100% internal overhead to handle control protocols over and above the datapath requirements. This 2X increase is an example of the dramatic I/O performance increase available in Virtex-II Pro Platform FPGAs.

Aurora Boosts FPGA-to-FPGA Links

Xilinx has developed a software macro called the Aurora core that provides easy 16-bit and 32-bit interfaces from FPGA-to-FPGA using one or more Rocket I/O MGTs. The Aurora core handles the framing, synchronization, and channel bonding tasks, allowing you to focus more on their application. A single MGT can provide a 16-bit or 32-bit FPGA-to-FPGA interface. The software Aurora core coupled with hard core (implemented in silicon) MGTs can create powerful serial-to-parallel and parallel-to-serial tranceivers.

Alternatively, multiple MGTs can be bonded together to form a higher-bandwidth interface. The Aurora macro ensures that channel-bonded data will appear on the same clock cycle at the other end of the communications link. The 8B/10B encoding method is used for the Aurora soft macro, giving an effective bandwidth of 10 gigabits per second for a set of four channel-bonded MGTs between two FPGAs.

Pre-Emphasis Goes the Distance

Another important feature of the Rocket I/O solution is pre-emphasis, which compensates for the filtering effects of FR-4 PC board material at gigabit speeds. Pre-emphasis boosts the output levels to compensate for the filtering effects of extended PCB traces. With pre-emphasis, PCB runs of 20 inches or longer can be supported reliably at speeds of 3.125 gigabits per second.

Figure 4 shows how two MGTs on separate FPGAs can be easily linked up to 20 inches apart using standard PC board transmission-line traces. Pre-emphasis improves the noise and jitter performance at these high speeds for longer PC traces where dispersion affects timing and voltage margins.

You can clearly see in Figure 5, which shows oscilloscope traces at the receiver without pre-emphasis and with 30% pre-emphasis. The pre-emphasis enables reliable signal transmission at 3.125 Gigabits per second over long PC board traces in standard FR-4 PC board material. Transmitter and receiver termination is provided internally to each MGT, eliminating the need for external termination of these transmission lines. The termination impedance can be set to 50 ohms or to 75 ohms.

No Pre-Emphasis

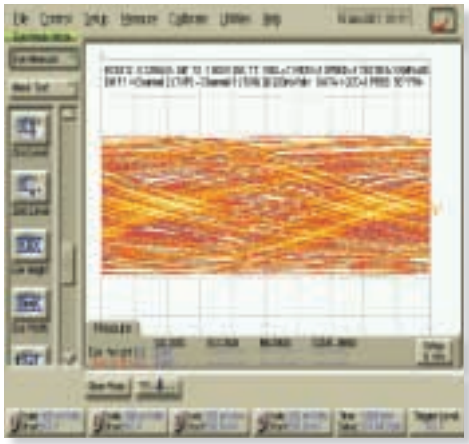

30% Pre-Emphasis

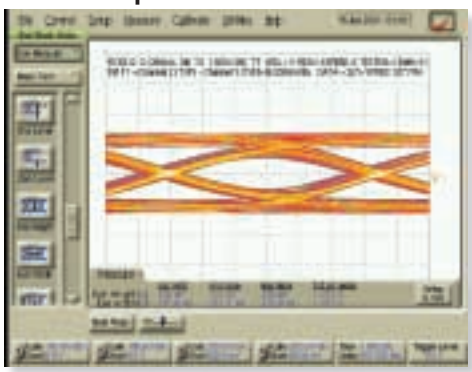

Figure 5 -Received signals after 50 inches FR-4 PC board material with no pre-emphasis and with 30% pre-emphasis in the Rocket I/O MGT.

Conclusion

The Rocket I/O MGTs enable high-speed interfaces for Virtex-II Pro Platform FPGAs. Standards such as InfiniBand Architecture, Fibre Channel, XAUI, and Gigabit Ethernet are directly supported. Inter-FPGA communications are greatly enhanced by Rocket I/O modules and the Aurora soft macro, which enable simple FPGA links with channel-bonded performance of over 10 gigabits per second. Using pre-emphasis, PCB trace runs of 20 inches or greater are possible at speeds of up to 3.125 gigabits per second per link. In the larger Virtex-II Pro devices, this represents a factor of two increase in total I/O performance available per FPGA. Clearly, Virtex-II Pro Platform FPGAs with Rocket I/O MGTs give you a competitive advantage, both in terms of performance and time to market.

New CoolRunner-II CPLD Development Kit

An ideal kit for applying CPLD technology to high-performance portable and battery-powered systems.

by Xilinx Staff

Insight Electronics recently introduced a complete solution for developing designs and applications based on the new high-performance, ultra-low power Xilinx CoolRunner™-II CPLD family. The Insight CoolRunner-II Development Kit enables experimentation with CoolRunner-II design concepts, including the investigation of multiple I/O standards, clock management, security, and CPLD power consumption.

> "Xilinx is pleased to have Insight providing timely support for our breakthrough CoolRunner-II RealDigital CPLDs. Through this development kit, users can evaluate the unbeatable combination of high performance and ultra low power that CoolRunner-II devices offer."
>
> – Steve Sharp, senior manager of Silicon Solutions Marketing.

"Releasing our CoolRunner-II Development Kit in tandem with the Xilinx CoolRunner-II family introduction continues Insight's successful development kit introduction strategy," said Jim Beneke, director of Technical Marketing at Insight. "This approach gives designers immediate access to the devices' full range of features, allowing them to verify the ultra-low power and high performance of the Xilinx second generation Fast Zero Power™ (FZP) technology as used in the CoolRunner-II family."

The Insight kit includes the 64 macrocell, 1.8V CoolRunner-II device; multiple options for power supplies, I/O voltages, and clock sources; a JTAG port; 45 user I/Os; a prototyping area; a two-digit LCD display; and two push-button switches. The Xilinx ISE WebPACK™ software contains all the development software you need, and can be downloaded at no cost from the Xilinx website, or is included in the WebPACK version of the kit.

"Xilinx is pleased to have Insight providing timely support for our breakthrough CoolRunner-II RealDigital CPLDs," stated Steve Sharp, senior manager of Silicon Solutions Marketing. "Through this development kit, users can evaluate the unbeatable combination of high performance and ultra low power that CoolRunner-II devices offer."

Price and Availability

Available for order now, Insight's CoolRunner-II Development Kit is priced at $95. For more information on this and other Xilinx design kits available from Insight, call 888-488-4133, or go to *www.insight-electronics.com/coolrunner2.*

About Insight Electronics

Insight Electronics is a member of The Memec Group, the world's leading distributor of proprietary advanced-technology semiconductors. With 52 divisions in the U.S., Mexico, Canada, and South America, Insight is the largest specialty semiconductor distributor in North and South America. Insight focuses its line card on leading semiconductor companies, enabling their sales team and field application engineers to devote their time and resources to mastery of supplier products, processes, and technologies. The company's value-added services, design expertise, inventory management, logistics, and e-business tools ensure that customers receive the products they need, when and where they need them. For more information, visit the Insight Website at *www.insight-electronics.com.*

Insight MEMEC

Real Performance.
Real Low-Power.
The RealDigital™ CPLD.

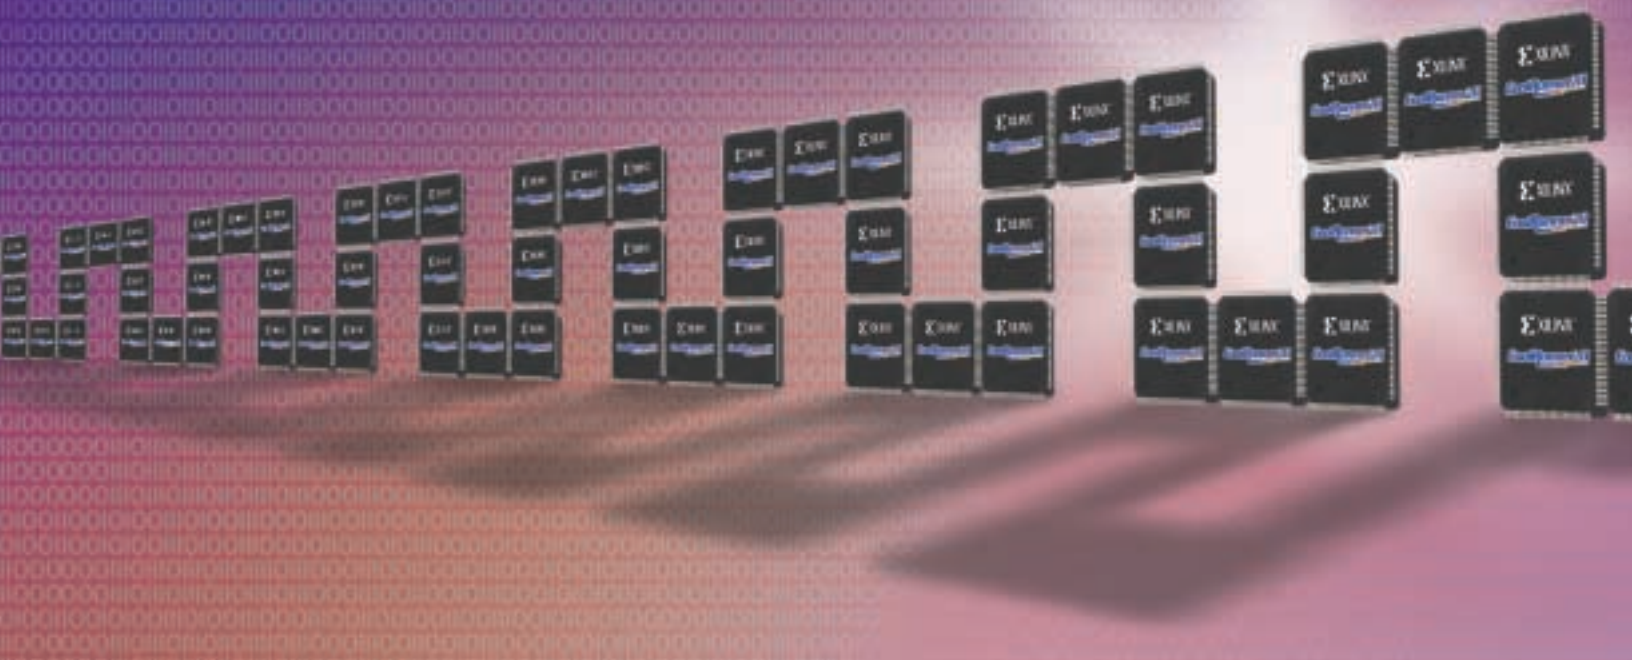

CoolRunner™-II The RealDigital CPLD is the breakthrough in high-performance and low-power you've been looking for. Shipping right now, the new CoolRunner™-II family of devices are the only CPLDs offering a 100% digital core, eliminating power-hungry analog sense amps.

A real-world feature set for real-world challenges

With complete I/O support including HSTL and SSTL, you've got the real-world interfacing you need today. Our unique Clock Divider means no more dividing the clock off-chip. And the Clock Doubler gives you speeds up to 400 MHz. The new RealDigital CPLDs also provide advanced Design Security. A comparison of 1.8V CPLD devices shows CoolRunner-II is still up to 20 times less power than our nearest competitor!

1.8V CORE VOLTAGE CPLD COMPETITIVE CHART

Manufacturer	Xilinx	Lattice	Altera
Device Family	**CoolRunner-II**	ispMACH4000C	None
Standby Current	**<100 µA**	2 mA	N/A
Clock Divider	**YES**	NO	N/A
Clock Doubler	**YES**	NO	N/A
I/O Standards Support	**LVTTL, LVCMOS, HSTL, SSTL**	LVTTL, LVCMOS	N/A
I/O Banks (max)	**4**	2	N/A

Only CoolRunner-II — with Fast Zero Power™ technology — provides ultra-high performance and ultra-low standby current. It's fast. It's cool. It's available now!

FREE downloadable software

The new CoolRunner-II RealDigital CPLDs are fully supported by the easy-to-use ISE WebPACK™, downloadable FREE via the Internet. Or choose our ISE 4.1i software — the fastest, most advanced development system in the industry, and part of the total solution offered only by Xilinx.

Find out more about the RealDigital CPLD, plus get your free CoolRunner-II resource CD by visiting **www.xilinx.com/digital** today.

XILINX®
The Programmable Logic Company℠

www.xilinx.com/digital

FORTUNE® 2002
100 BEST COMPANIES TO WORK FOR

Virtex-II Pro Platform FPGAs: Introduction and Overview

XCELL (v1.0) January 11, 2002

Advance Product Specification

Summary of Virtex-II Pro Features

- High-performance Platform FPGA solution including
  - Up to sixteen Rocket I/O™ embedded multi-gigabit transceiver blocks (based on Mindspeed's SkyRail™ technology)
  - Up to four IBM® PowerPC™ RISC processor blocks
- Based on Virtex™-II Platform FPGA technology
  - Flexible logic resources
  - SRAM-based in-system configuration
  - Active Interconnect™ technology
  - SelectRAM™ memory hierarchy
  - Dedicated 18-bit x 18-bit multiplier blocks
  - High-performance clock management circuitry
  - SelectI/O™-Ultra technology
  - Digitally Controlled Impedance (DCI) I/O

The members and resources of the Virtex-II Pro family are shown in Table 1.

Rocket I/O™ Features

- Full-duplex serial transceiver (SERDES) capable of baud rates from 622 Mb/s to 3.125 Gb/s
- 80 Gb/s duplex data rate (16 channels)
- Monolithic clock synthesis and clock recovery (CDR)
- Fibre Channel, Gigabit Ethernet, 10-Gbit Attachment Unit Interface (XAUI), and Infiniband-compliant transceivers
- 8-, 16-, or 32-bit selectable internal FPGA interface
- 8B/10B encoder and decoder
- 50Ω / 75Ω on-chip selectable transmit and receive terminations
- Programmable comma detection
- Channel bonding support (two to sixteen channels)
- Rate matching via insertion/deletion characters
- Four levels of selectable pre-emphasis
- Five levels of output differential voltage
- Per-channel internal loopback modes
- 2.5V transceiver supply voltage

PowerPC RISC Core Features

- Embedded 300+ MHz Harvard architecture core
- Low power consumption: 0.9 mW/MHz
- Five-stage data path pipeline
- Hardware multiply/divide unit
- Thirty-two 32-bit general purpose registers
- 16 KB two-way set-associative instruction cache
- 16 KB two-way set-associative data cache
- Memory Management Unit (MMU)
  - 64-entry unified Translation Look-aside Buffers (TLB)
  - Variable page sizes (1 KB to 16 MB)
- Dedicated on-chip memory (OCM) interface
- Supports IBM CoreConnect™ bus architecture
- Debug and trace support
- Timer facilities

Table 1: **Virtex-II Pro Field-Programmable Gate Array Family Members**

Device	Rocket I/O Transceiver Blocks	PowerPC Processor Blocks	CLB (1 CLB = 4 slices = Max 128 bits)			18 X 18 Bit Multiplier Blocks	Block SelectRAM		DCMs	Max I/O Pads
			Array Row x Col	Slices	Maximum Distributed RAM Kbits		18 Kbit Blocks	Max Block RAM Kbits		
XC2VP2	4	0	16 x 22	1,408	44	12	12	216	4	204
XC2VP4	4	1	40 x 22	3,008	94	28	28	504	4	348
XC2VP7	8	1	40 x 34	4,928	154	44	44	792	4	396
XC2VP20	8	2	56 x 46	9,280	290	88	88	1,584	8	564
XC2VP50	16	4	88 x 70	22,592	706	216	216	3,888	8	852

Virtex-II Pro Platform FPGA Technology

- SelectRAM memory hierarchy
  - Up to 4 Mbit of True Dual-Port RAM in 18-Kb block SelectRAM resources
  - Up to 706 Kb of distributed SelectRAM resources
  - High-performance interfaces to external memory
- Arithmetic functions
  - Dedicated 18-bit x 18-bit multiplier blocks
  - Fast look-ahead carry logic chains
- Flexible logic resources
  - Up to 45,184 internal registers / latches with Clock Enable
  - Up to 45,184 Look-Up Tables (LUTs) or cascadable variable (1 to 16 bits) shift registers
  - Wide multiplexers and wide-input function support
  - Horizontal cascade chain and Sum-of-Products support
  - Internal 3-state busing
- High-performance clock management circuitry
  - Up to eight Digital Clock Manager (DCM) modules
    - Precise clock de-skew
    - Flexible frequency synthesis
    - High-resolution phase shifting
  - 16 global clock multiplexer buffers in all parts
- Active Interconnect™ technology
  - Fourth-generation segmented routing structure
  - Fast, predictable routing delay, independent of fanout
  - Deep sub-micron noise immunity benefits
- SelectI/O-Ultra™ technology
  - Up to 852 user I/Os
  - Twenty two single-ended standards and five differential standards
  - Programmable LVTTL and LVCMOS sink/source current (2 mA to 24 mA) per I/O
  - Digitally Controlled Impedance (DCI) I/O: on-chip termination resistors for single-ended I/O standards
  - PCI support (designated banks only)
  - Differential signaling
    - 840 Mb/s Low-Voltage Differential Signaling I/O (LVDS) with current mode drivers
    - Bus LVDS I/O
    - HyperTransport (LDT) I/O with current driver buffers
    - Built-in DDR input and output registers
  - Proprietary high-performance SelectLink technology for communications between Xilinx devices
    - High-bandwidth data path
    - Double Data Rate (DDR) link
    - Web-based HDL generation methodology
- SRAM-based in-system configuration
  - Fast SelectMAP™ configuration
  - Triple Data Encryption Standard (DES) security option (bitstream encryption)
  - IEEE1532 support
  - Partial reconfiguration
  - Unlimited reprogrammability
  - Readback capability
- Supported by Xilinx Foundation™ and Alliance™ series development systems
  - Integrated VHDL and Verilog design flows
  - ChipScope™ Integrated Logic Analyzer
- 0.13-µm, nine-layer copper process with 90 nm high-speed transistors
- 1.5V (VCCINT) core power supply, dedicated 2.5V VCCAUX auxiliary and V_{CCO} I/O power supplies
- IEEE 1149.1 compatible boundary-scan logic support
- Flip-Chip and Wire-Bond Ball Grid Array (BGA) packages in standard 1.00 mm pitch
- Each device 100% factory tested

General Description

The Virtex-II Pro family is a platform FPGA for designs that are based on IP cores and customized modules. The family incorporates multi-gigabit transceivers and PowerPC CPU cores in Virtex-II Pro series FPGA architecture. It empowers complete solutions for telecommunication, wireless, networking, video, and DSP applications.

The leading-edge 0.13µm CMOS nine-layer copper process and the Virtex-II Pro architecture are optimized for high performance designs in a wide range of densities. Combining a wide variety of flexible features and IP cores, the Virtex-II Pro family enhances programmable logic design capabilities and is a powerful alternative to mask-programmed gates arrays.

Architecture

Virtex-II Pro Array Overview

Virtex-II Pro devices are user-programmable gate arrays with various configurable elements and embedded cores optimized for high-density and high-performance system designs. Virtex-II Pro implements the following functionality:

- Embedded high-speed serial transceivers enable data bit rate up to 3.125 Gb/s per channel.
- Embedded IBM PowerPC 405 RISC CPU cores provide performance of 300+ MHz.
- SelectI/O-Ultra blocks provide the interface between package pins and the internal configurable logic. Most popular and leading-edge I/O standards are supported by the programmable IOBs.
- Configurable Logic Blocks (CLBs) provide functional elements for combinatorial and synchronous logic, including basic storage elements. BUFTs (3-state buffers) associated with each CLB element drive dedicated segmentable horizontal routing resources.
- Block SelectRAM memory modules provide large 18-Kb storage elements of True Dual-Port RAM.
- Embedded multiplier blocks are 18-bit x 18-bit dedicated multipliers.
- DCM (Digital Clock Manager) blocks provide self-calibrating, fully digital solutions for clock distribution delay compensation, clock multiplication and division, and coarse- and fine-grained clock phase shifting.

A new generation of programmable routing resources called Active Interconnect Technology interconnects all of these elements. The general routing matrix (GRM) is an array of routing switches. Each programmable element is tied to a switch matrix, allowing multiple connections to the general routing matrix. The overall programmable interconnection is hierarchical and designed to support high-speed designs.

All programmable elements, including the routing resources, are controlled by values stored in static memory cells. These values are loaded in the memory cells during configuration and can be reloaded to change the functions of the programmable elements.

Virtex-II Pro Features

This section briefly describes Virtex-II Pro features.

Rocket I/O Multi-Gigabit Transceiver Cores

The Rocket I/O Multi-Gigabit Transceiver core, based on Mindspeed's SkyRail technology, is a flexible parallel-to-serial and serial-to-parallel transceiver embedded core used for high-bandwidth interconnection between buses, backplanes, or other subsystems.

Multiple user instantiations in an FPGA are possible, providing up to 80 Gb/s of full duplex raw data transfer. Each channel can be operated at a maximum data transfer rate of 3.125 Gb/s.

Each Rocket I/O core implements the following functionality:

- Serializer and deserializer (SERDES)
- Monolithic clock synthesis and clock recovery (CDR)
- Fibre Channel, Gigabit Ethernet, XAUI, and Infiniband compliant transceivers
- 8-, 16-, or 32-bit selectable FPGA interface
- 8B/10B encoder and decoder with bypassing option on each channel
- Channel bonding support (two to sixteen channels)
  - Elastic buffers for inter-chip deskewing and channel-to-channel alignment
- Receiver clock recovery tolerance of up to 75 non-transitioning bits
- 50Ω / 75Ω on-chip selectable TX and RX terminations
- Programmable comma detection
- Rate matching via insertion/deletion characters
- Automatic lock-to-reference function
- Optional TX and RX data inversion
- Four levels of pre-emphasis support
- Per-channel serial and parallel transmitter-to-receiver internal loopback modes
- Cyclic Redundancy Check (CRC) support

PowerPC 405

The PPC405 RISC CPU can execute instructions at a sustained rate of one instruction per cycle. On-chip instruction and data cache reduce design complexity and improve system throughput.

PPC405 features include:

- PowerPC RISC CPU
  - Implements the PowerPC User Instruction Set Architecture (UISA) and extensions for embedded applications
  - Thirty-two 32-bit general purpose registers (GPRs)
  - Static branch prediction
  - Five-stage pipeline with single-cycle execution of most instructions, including loads/stores
  - Unaligned and aligned load/store support to cache, main memory, and on-chip memory
  - Hardware multiply/divide for faster integer arithmetic (4-cycle multiply, 35-cycle divide)
  - Enhanced string and multiple-word handling
  - Big/little endian operation support
- Storage Control
  - Separate instruction and data cache units, both two-way set-associative and non-blocking
  - Eight words (32 bytes) per cache line
  - 16 KB array Instruction Cache Unit (ICU), 16 KB array Data Cache Unit (DCU)

  - Operand forwarding during instruction cache line fill
  - Copy-back or write-through DCU strategy
  - Doubleword instruction fetch from cache improves branch latency
- Virtual mode memory management unit (MMU)
  - Translation of the 4 GB logical address space into physical addresses
  - Software control of page replacement strategy
  - Supports multiple simultaneous page sizes ranging from 1 KB to 16 MB
- OCM controllers provide dedicated interfaces between Block SelectRAM memory and processor core instruction and data paths for high-speed access
- PowerPC timer facilities
  - 64-bit time base
  - Programmable interval timer (PIT)
  - Fixed interval timer (FIT)
  - Watchdog timer (WDT)
- Debug Support
  - Internal debug mode
  - External debug mode
  - Debug Wait mode
  - Real Time Trace debug mode
  - Enhanced debug support with logical operators
  - Instruction trace and trace-back support
  - Forward or backward trace
- Two hardware interrupt levels support
- Advanced power management support

Input/Output Blocks (IOBs)

IOBs are programmable and can be categorized as follows:

- Input block with an optional single data rate (SDR) or double data rate (DDR) register
- Output block with an optional SDR or DDR register and an optional 3-state buffer to be driven directly or through SDR or DDR register
- Bi-directional block (any combination of input and output configurations)

These registers are either edge-triggered D-type flip-flops or level-sensitive latches.

IOBs support the following single-ended I/O standards:

- LVTTL
- LVCMOS (3.3V, 2.5V, 1.8V, and 1.5V)
- PCI (33 and 66 MHz)
- GTL and GTLP
- HSTL 1.5V and 1.8V (Class I, II, III, and IV)
- SSTL 3.3V and 2.5V, (Class I and II)

The digitally controlled impedance (DCI) I/O feature automatically provides on-chip termination for each single-ended I/O standard.

The IOB elements also support the following differential signaling I/O standards:

- LVDS and Extended LVDS (2.5V only)
- BLVDS (Bus LVDS)
- ULVDS
- LDT

Two adjacent pads are used for each differential pair. Two or four IOB blocks connect to one switch matrix to access the routing resources.

Configurable Logic Blocks (CLBs)

CLB resources include four slices and two 3-state buffers. Each slice is equivalent and contains:

- Two function generators (F & G)
- Two storage elements
- Arithmetic logic gates
- Large multiplexers
- Wide function capability
- Fast carry look-ahead chain
- Horizontal cascade chain (OR gate)

The function generators F & G are configurable as 4-input look-up tables (LUTs), as 16-bit shift registers, or as 16-bit distributed SelectRAM memory.

In addition, the two storage elements are either edge-triggered D-type flip-flops or level-sensitive latches.

Each CLB has internal fast interconnect and connects to a switch matrix to access general routing resources.

Block SelectRAM Memory

The block SelectRAM memory resources are 18 Kb of True Dual-Port RAM, programmable from 16K x 1 bit to 512 x 36 bit, in various depth and width configurations. Each port is totally synchronous and independent, offering three “read-during-write” modes. Block SelectRAM memory is cascadable to implement large embedded storage blocks. Supported memory configurations for dual-port and single-port modes are shown in Table 2.

Table 2: **Dual-Port and Single-Port Configurations**

16K x 1 bit	4K x 4 bits	1K x 18 bits
8K x 2 bits	2K x 9 bits	512 x 36 bits

18 X 18-Bit Multipliers

A multiplier block is associated with each SelectRAM memory block. The multiplier block is a dedicated 18 x 18-bit 2s complement signed multiplier, and is optimized for operations based on the block SelectRAM content on one port. The 18 x 18 multiplier can be used independently of the block SelectRAM resource. Read/multiply/accumulate operations and DSP filter structures are extremely efficient.

Both the SelectRAM memory and the multiplier resource are connected to four switch matrices to access the general routing resources.

Global Clocking

The DCM and global clock multiplexer buffers provide a complete solution for designing high-speed clock schemes.

Up to eight DCM blocks are available. To generate deskewed internal or external clocks, each DCM can be used to eliminate clock distribution delay. The DCM also provides 90, 180, and 270-degree phase-shifted versions of its output clocks. Fine-grained phase shifting offers high-resolution phase adjustments in increments of 1/256 of the clock period. Very flexible frequency synthesis provides a clock output frequency equal to a fractional or integer multiple of the input clock frequency. For exact timing parameters, see *Virtex-II Pro Platform FPGAs: DC and Switching Characteristics*.

Virtex-II Pro devices have 16 global clock MUX buffers, with up to eight clock nets per quadrant. Each clock MUX buffer can select one of the two clock inputs and switch glitch-free from one clock to the other. Each DCM can send up to four of its clock outputs to global clock buffers on the same edge. Any global clock pin can drive any DCM on the same edge.

Routing Resources

The IOB, CLB, block SelectRAM, multiplier, and DCM elements all use the same interconnect scheme and the same access to the global routing matrix. Timing models are shared, greatly improving the predictability of the performance of high-speed designs.

There are a total of 16 global clock lines, with eight available per quadrant. In addition, 24 vertical and horizontal long lines per row or column, as well as massive secondary and local routing resources, provide fast interconnect. Virtex-II Pro buffered interconnects are relatively unaffected by net fanout, and the interconnect layout is designed to minimize crosstalk.

Horizontal and vertical routing resources for each row or column include:

- 24 long lines
- 120 hex lines
- 40 double lines
- 16 direct connect lines (total in all four directions)

Boundary Scan

Boundary-scan instructions and associated data registers support a standard methodology for accessing and configuring Virtex-II Pro devices, complying with IEEE standards 1149.1 and 1532. A system mode and a test mode are implemented. In system mode, a Virtex-II Pro device will continue to function while executing non-test boundary-scan instructions. In test mode, boundary-scan test instructions control the I/O pins for testing purposes. The Virtex-II Pro Test Access Port (TAP) supports BYPASS, PRELOAD, SAMPLE, IDCODE, and USERCODE non-test instructions. The EXTEST, INTEST, and HIGHZ test instructions are also supported.

Configuration

Virtex-II Pro devices are configured by loading the bitstream into internal configuration memory using one of the following modes:

- Slave-serial mode
- Master-serial mode
- Slave SelectMAP mode
- Master SelectMAP mode
- Boundary-Scan mode (IEEE 1532)

A Data Encryption Standard (DES) decryptor is available on-chip to secure the bitstreams. One or two triple-DES key sets can be used to optionally encrypt the configuration data.

The Xilinx System Advanced Configuration Enviornment (System ACE) family offers high-capacity and flexible solution for FPGA configuration as well as program/data storage for the processor. See **DS080**, ***System ACE CompactFlash Solution*** for more information.

Readback and Integrated Logic Analyzer

Configuration data stored in Virtex-II Pro configuration memory can be read back for verification. Along with the configuration data, the contents of all flip-flops/latches, distributed SelectRAM, and block SelectRAM memory resources can be read back. This capability is useful for real-time debugging.

The Xilinx ChipScope Pro Integrated Logic Analyzer (ILA) cores and Integrated Bus Analyzer (IBA) cores, along with the ChipScope Pro Analyzer software, provide a complete solution for accessing and verifying user designs within Virtex-II Pro devices.

IP Core and Reference Support

Intellectual Property is part of the Platform FPGA solution. In addition to the existing FPGA fabric cores, the list below shows some of the currently available hardware and software intellectual properties specially developed for Virtex-II Pro by Xilinx. Each IP core is modular, portable, Real-Time Operating System (RTOS) independent, and CoreConnect compatible for ease of design migration. Refer to **www.xilinx.com** for the latest and most complete list of cores.

Hardware Cores

- Bus Infastructure cores (arbiters, bridges, and more)
- Memory cores (Flash, SRAM, and more)
- Peripheral cores (UART, IIC, and more)
- Networking cores (ATM, Ethernet, and more)

Software Cores

- Boot code
- Test code
- Device drivers
- Protocol stacks
- RTOS integration
- Customized board support package

Virtex-II Pro Device/Package Combinations and Maximum I/Os

Offerings include ball grid array (BGA) packages with 1.0 mm pitch. In addition to traditional wire-bond interconnects, flip-chip interconnect is used in some of the BGA offerings. The use of flip-chip interconnect offers more I/Os than is possible in wire-bond versions of the similar packages. Flip-chip construction offers the combination of high pin count and excellent power dissipation.

The Virtex-II Pro device/package combination table (Table 3) details the maximum number of I/Os for each device and package using wire-bond or flip-chip technology.

- FG denotes wire-bond fine-pitch BGA (1.00 mm pitch).
- FF denotes flip-chip fine-pitch BGA (1.00 mm pitch).
- BF denotes flip-chip fine-ptich BGA (1.27 mm pitch).

Table 3: **Virtex-II Pro Device/Package Combinations and Maximum Number of Available I/Os (Advance Information)**

Package	Pitch (mm)	Size (mm)	User Available I/Os				
			XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP50
FG256	1.00	17 x 17	140	140			
FG456	1.00	23 x 23	156	248	248		
FF672	1.00	27 x 27	204	348	396		
FF896	1.00	31 x 31			396	556	
FF1152	1.00	35 x 35				564	692
FF1517	1.00	40 x 40					852
BF957	1.27	40 x 40				564	584

For digital video, Spartan®-IIE is the clear choice.

The new Spartan-IIE FPGAs deliver over 20 billion MACs/s, giving you the DSP performance required for high-resolution digital images. Together with advanced signaling standards including LVDS, HSTL and SSTL, you've got the speed you need for digital video processing.

From Picture to Pixel, the Complete Solution

Through capture, process and display, Spartan-IIE FPGAs bring the total solution into focus. Encryption IP, operating at high speed video rates, offers a new level of security protecting every pixel. And the desktop programmability of Spartan-IIE FPGAs solves your time-to-market pressures as design cycles get shorter.

The eSP Web Portal . . . a Unique Designer's Resource

The Xilinx eSP web portal gives you all the support you could wish for as standards continue to rapidly change. Simply log on to **xilinx.com/esp** and you'll get immediate access to reference boards, system block diagrams, a full spectrum of technology & standards tutorials, and much more. You'll soon see why eSP is one of the most successful websites of its kind.

Visit ***www.xilinx.com/spartan2e/dvt*** today and find out why Spartan-IIE FPGAs are the clear choice for digital video design.

XILINX®

The Programmable Logic Company℠

www.xilinx.com/spartan2e/dvt

FORTUNE® 2001 100 BEST COMPANIES TO WORK FOR

Xilinx FPGA Product Selection Matrix

Device	Rocket I/O™ Transceiver Blocks	PowerPC™ Processor Blocks	CLB (1 CLB = 4 slices = Max 128 bits) Array Row x Col	CLB Slices	Maximum Distributed RAM Kbits	18 x 18 Bit Multiplier Blocks	Block SelectRAM 18 Kbit Blocks	Block SelectRAM Max Block RAM Kbits	DCMs	Max I/O Pads
XC2VP2	4	0	16 x 22	1,408	44	12	12	216	4	204
XC2VP4	4	1	40 x 22	3,008	94	28	28	504	4	348
XC2VP7	8	1	40 x 34	4,928	154	44	44	792	4	396
XC2VP20	8	2	56 x 46	9,280	290	88	88	1,584	8	564
XC2VP50	16	4	88 x 70	22,592	706	216	216	3,888	8	852

	CLB Resources						BLK RAM		CLK Resources					I/O Features				Speed				
	System Gates (see note 1)	CLB Array (Row X Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM Blocks	Block RAM Bits	# Dedicated Multipliers	DLL Frequency (min/max)	# DLL's	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family		Config. Memory (Bits)
Virtex-II Family — 1.5 Volt																		**.15um Eight Layer Metal Process**				
XC2V40	40K	8 x 8	256	576	512	8K	4	72K	4	24/420	4	DCM	DCM	YES	44	88	LDT-25, LVPECL-33,	-4 -5 -6	-4 -5	ISP	OTP	0.4M
XC2V80	80K	16 x8	512	1,152	1,024	16K	8	144K	8	24/420	4	DCM	DCM	YES	60	120	LVDS-33, LVDS-25,	-4 -5 -6	-4 -5			0.6M
XC2V250	250K	24 x16	1,536	3,456	3,072	48K	24	432K	24	24/420	8	DCM	DCM	YES	100	200	LVDSEXT-33, LVDSEXT-25,	-4 -5 -6	-4 -5			1.7M
XC2V500	500K	32 x 24	3,072	6,912	6,144	96K	32	576K	32	24/420	8	DCM	DCM	YES	132	264	BLVDS-25, ULVDS-25,	-4 -5 -6	-4 -5			2.8M
XC2V1000	1M	40 x 32	5,120	11,520	10,240	160K	40	720K	40	24/420	8	DCM	DCM	YES	216	432	LVTTL, LVCMOS33,	-4 -5 -6	-4 -5			4.1M
XC2V1500	1.5M	48 x 40	7,680	17,280	15,360	240K	48	864K	48	24/420	8	DCM	DCM	YES	264	528	LVCMOS25, LVCMOS18,	-4 -5 -6	-4 -5			5.7M
XC2V2000	2M	56 x 48	10,752	24,192	21,504	336K	56	1008K	56	24/420	8	DCM	DCM	YES	312	624	LVCMOS15, PCI33, PCI66,	-4 -5 -6	-4 -5			7.5M
XC2V3000	3M	64 x 56	14,336	32,256	28,672	448K	96	1728K	96	24/420	12	DCM	DCM	YES	360	720	PCI-X, GTL, GTL+, HSTL I,	-4 -5 -6	-4 -5			10.5M
XC2V4000	4M	80 x 72	23,040	51,840	46,080	720K	120	2160K	120	24/420	12	DCM	DCM	YES	456	912	HSTL II, HSTL III, HSTL IV,	-4 -5 -6	-4 -5			15.7M
XC2V6000	6M	96 x 88	33,792	76,032	67,584	1056K	144	2592K	144	24/420	12	DCM	DCM	YES	552	1104	SSTL21, SSTL211,	-4 -5 -6	-4 -5			21.9M
XC2V8000	8M	112 x 104	46,592	104,832	93,184	1456K	168	3024K	168	24/420	12	DCM	DCM	YES	554	1108	SSTL3 I, SSTL3 II	-4 -5	-4			29.1M
Virtex-E Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XCV50E	72K	16 x 24	768	1,728	1,536	24K	16	64K	NA	25/350	8	YES	YES	NA	88	176		-6 -7 -8	-6 -7	ISP	OTP	0.6M
XCV100E	128K	20 x 30	1,200	2,700	2,400	37.5K	20	80K	NA	25/350	8	YES	YES	NA	98	196		-6 -7 -8	-6 -7			0.9M
XCV200E	306K	28 x 42	2,352	5,292	4,704	73.5K	28	112K	NA	25/350	8	YES	YES	NA	142	284	LVTTL, LVCMOS2,	-6 -7 -8	-6 -7			1.45M
XCV300E	412K	32 x 48	3,072	6,912	6,144	96K	32	128K	NA	25/350	8	YES	YES	NA	158	316	LVCMOS18, PCI33,	-6 -7 -8	-6 -7			1.88M
XCV400E	570K	40 x 60	4,800	10,800	9,600	150K	40	160K	NA	25/350	8	YES	YES	NA	202	404	PCI66, GTL, GTL+,	-6 -7 -8	-6 -7			2.7M
XCV600E	986K	48 x 72	6,912	15,552	13,824	216K	72	288K	NA	25/350	8	YES	YES	NA	256	512	HSTL I, HSTL III, HSTL IV,	-6 -7 -8	-6 -7			3.97M
XCV1000E	1,569K	64 x 96	12,288	27,648	24,576	384K	96	384K	NA	25/350	8	YES	YES	NA	330	660	SSTL3 I, SSTL3 II,	-6 -7 -8	-6 -7			6.6M
XCV1600E	2,188K	72 x 180	25,920	34,992	51,840	486K	144	576K	NA	25/350	8	YES	YES	NA	362	724	SSTL21, SSTL211, BLVDS,	-6 -7 -8	-6 -7			8.4M
XCV2000E	2,542K	80 x 120	19,200	43,200	38,400	600K	160	640K	NA	25/350	8	YES	YES	NA	402	804	LVDS, LVPECL	-6 -7 -8	-6 -7			10.2M
XCV2600E	3,264K	92 x 138	25,392	57,132	50,784	793.5K	184	736K	NA	25/350	8	YES	YES	NA	402	804		-6 -7 -8	-6 -7			13M
XCV3200E	4,074K	104 x 156	32,448	73,008	64,896	1014K	208	832K	NA	25/350	8	YES	YES	NA	402	804		-6 -7 -8	-6 -7			16.3M
Virtex-EM Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XCV405E	1.31M	40 x 60	4,800	10,800	9,600	150K	140	560K	NA	25/350	8	YES	YES	NA	202	404	Same As	-6 -7 -8	-6 -7	ISP	OTP	3.43M
XCV812E	2.54M	56 x 84	9,408	21,168	18,816	294K	280	1120K	NA	25/350	8	YES	YES	NA	278	556	Virtex-E	-6 -7 -8	-6 -7			6.52M
Spartan-IIE Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/320	4	YES	YES	NA	84	182	LVTTL,LVCMOS2, LVCMOS18,	-6 -7	-6	ISP	OTP	0.6M
XC2S100E	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/320	4	YES	YES	NA	86	202	PCI33, PCI66, GTL, GTL+,	-6 -7	-6			0.9M
XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/320	4	YES	YES	NA	114	263	HSTL I, HSTL III, HSTL IV, SSTL3 I,	-6 -7	-6			1.1M
XC2S200E	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/320	4	YES	YES	NA	120	289	SSTL3 II, SSTL2 I,SSTL2 II, AGP-2X,	-6 -7	-6			1.4M
XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	NA	25/320	4	YES	YES	NA	120	329	CTT, LVDS, BLVDS, LVPECL	-6 -7	-6			1.9M
Spartan-II Family — 2.5 Volt																		**.22/.18um Six Layer Metal Process**				
XC2S15	15K	8 x 12	192	432	384	6K	4	16K	NA	25/200	4	YES	YES	NA	NA	86	LVTTL, LVCMOS2,	-5 -6	-5	ISP	OTP	0.2M
XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	NA	25/200	4	YES	YES	NA	NA	132	PCI33 (3.3V & 5V),	-5 -6	-5			0.4M
XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/200	4	YES	YES	NA	NA	176	PCI66 (3.3V), GTL, GTL+,	-5 -6	-5			0.6M
XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/200	4	YES	YES	NA	NA	196	HSTL I, HSTL III, HSTL IV,	-5 -6	-5			0.8M
XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/200	4	YES	YES	NA	NA	260	SSTL3 I, SSTL3 II, SSTL2 I,	-5 -6	-5			1.1M
XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/200	4	YES	YES	NA	NA	284	SSTL2 II, AGP-2X, CTT	-5 -6	-5			1.4M

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. A Logic Cell is defined as a 4 input LUT and a register
DCM – Digital Clock Management

Important: Verify all Data with Device Data Sheet

Xilinx FPGA Package Options and User I/O

VIRTEX-II PRO

Package	Pitch (mm)	Size (mm)	User Available I/Os: XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP50
FG256	1.00	17 x 17	140	140			
FG456	1.00	23 x 23	156	248	248		
FF672	1.00	27 x 27	204	348	396		
FF896	1.00	31 x 31			396	556	
FF1152	1.00	35 x 35				564	692
FF1517	1.00	40 x 40					852
BF957	1.27	40 x 40				564	584

			VIRTEX-II											VIRTEX-E												SPARTAN-IIE					SPARTAN-II					
			Virtex-II (1.5V)											Virtex-E (1.8V)										V-EM		Spartan-IIE (1.8V)					Spartan-II (2.5V)					
			XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000	XCV50E	XCV100E	XCV200E	XCV300E	XCV400E	XCV600E	XCV1000E	XCV1600E	XCV2000E	XCV2600E	XCV3200E	XCV405E	XCV812E	XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S15	XC2S30	XC2S50	XC2S100	XC2S150	XC2S200
Pins	Body Size	I/O's	88	120	200	264	432	528	624	720	912	1104	1296	176	176	284	316	404	512	660	724	804	804	804	404	556	182	202	263	289	329	86	132	176	196	260	284
PQFP Packages (PQ)																																					
208	28 x 28 mm																									146	146	146	146	146		132	140	140	140	140
240	32 x 32 mm													158	158	158	158	158	158	158																	
HQFP Packages (HQ)																																					
240	32 x 32 mm														158				158																		
VQFP Packages (VQ)																																					
100	14 x 14 mm																														60	60				
TQFP Packages (TQ)																																					
144	20 x 20 mm																									102	102				86	92	92	92		
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)																																					
144	12 x 12 mm		88	92	92									94	94	94															86	92				
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																																					
352	40 x 40 mm														196	260	260																				
432	40 x 40 mm																316	316	316																		
560	42.5 x 42.5 mm																	404	404	404	404	404			404	404											
575	31 x 31 mm						328	392	408																												
728	35 x 35 mm								456	516																											
FGA Packages (FT) — wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																																					
256	17 x 17 mm																									182	182	182	182	182						
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)																																					
256	17 x 17 mm		88	120	172	172	172							176	176	176	176																176	176	176	176
456	23 x 23 mm				200	264	324									284	312										202	263	289	329				196	260	284
676	27 x 27 mm							392	456	484								404	444						404											
680	40 x 40 mm																		512	512	512	512														
860	42.5 x 42.5 mm																			660	660	660														
900	31 x 31 mm																		512	660	700					556										
1156	35 x 35 mm																			660	724	804	804	804													
FFA Packages (FF) — flip-chip fine-pitch BGA (1.0 mm ball spacing)																																					
896	31 x 31 mm						432	528	624																												
1152	35 x 35 mm									720	824	824	824																								
1517	40 x 40 mm										912	1104	1108																								
BFA Packages (BF) — flip-chip fine-pitch BGA (1.27 mm ball spacing)																																					
957	40 x 40 mm								624	684	684	684	684																								

Note: Virtex-II packages FG456 and FG676 are footprint compatible.
Virtex-II packages FF896 and FF1152 are footprint compatible.
Important: Verify all Data with Device Data Sheet

Numbers indicated in the matrix are the maximum number of user I/O's for that package and device combination.

Xilinx CPLD Product Selection Matrix

PRODUCT SELECTION MATRIX

						I/O Features		Speed			Clocking	
	System Gates	Macrocells	Product terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	Max. I/O	I/O Banking	Min. Pin-to-Pin Logic Delay (ns)	Commercial Speed Grades (fastest to slowest)	Industrial Speed Grades (fastest to slowest)	Global Clocks	Product Term Clocks per Function Block
XC9500XV Family — 2.5 Volt												
XC9536XV	800	36	90	1.8/2.5/3.3	1.8/2.5/3.3	36	1	3.5	-3 -4 -5 -7	-7	3	18
XC9572XV	1600	72	90	1.8/2.5/3.3	1.8/2.5/3.3	72	1	4	-4 -5 -7	-7	3	18
XC95144XV	3200	144	90	1.8/2.5/3.3	1.8/2.5/3.3	117	2	4	-4 -5 -7	-7	3	18
XC95288XV	6400	288	90	1.8/2.5/3.3	1.8/2.5/3.3	192	4	5	-5 -7 -10	-10	3	18
XC9500XL Family — 3.3 Volt												
XC9536XL	800	36	90	2.5/3.3/5	2.5/3.3	36		5	-5 -7 -10	-7 -10	3	18
XC9572XL	1600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	3	18
XC95144XL	3200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	3	18
XC95288XL	6400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	3	18
CoolRunner XPLA3 — 3.3 Volt												
XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	4	16
XCR3064XL	1500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	4	16
XCR3128XL	3000	128	48	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	4	16
XCR3256XL	6000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	4	16
XCR3384XL	9000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	4	16
XCR3512XL	12000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	4	16
CoolRunner-II Family — 1.8 Volt												
XC2C32	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	1	3.5	-3 -4 -6	-4 -6	3	17
XC2C64	1500	64	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	64	1	4	-4 -5 -7	-5 -7	3	17
XC2C128	3000	128	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	100	2	4.5	-4 -6 -7	-6 -7	3	17
XC2C256	6000	256	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	184	2	5	-5 -6 -7	-6 -7	3	17
XC2C384	9000	384	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	240	4	6	-6 -7 -10	-7 -10	3	17
XC2C512	12000	512	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	270	4	6	-6 -7 -10	-7 -10	3	17

PACKAGE OPTIONS AND USER I/O

		XC9500XV				XC9500XL				CoolRunner XPLA3						CoolRunner II					
	Body Size	XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9536XL	XC9572XL	XC95144XL	XC95288XL	XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512
PLCC Packages (PC)																					
44	17.5 x 17.5 mm	34	34			34	34			36	36					33	33				
PQFP Packages (PQ)																					
208	28 x 28 mm				168				168				164	172	180				173	173	173
VQFP Packages (VQ)																					
44	12 x 12 mm	34	34			34	34			36	36					33	33				
64	12 x 12 mm					36	52														
100	16 x 16 mm										68	84					64	80	80		
TQFP Packages (TQ)																					
100	14 x 14 mm		72	81			72	81													
144	20 x 20 mm			117	117			117	117			108	120	118				100	118	118	
Chip Scale Packages (CP) — wire-bond chip-scale BGA (0.5 mm ball spacing)																					
56	6 X 6 mm										48					33	45				
Chip Scale Packages (CS) — wire-bond chip-scale BGA (0.8 mm ball spacing)																					
48	7 x 7mm	36	38			36	38			36	40										
144	12 x 12 mm			117				117				108									
280	16 x 16 mm				192				192				164								
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																					
256	27 x 27 mm								192										184	212	212
FBGA Packages (FG) — wire-bond Fineline BGA (1.0 mm ball spacing)																					
256	17 x 17 mm				192				192				164	212	212						
324	23 x 23 mm													220	260					240	270

Important: Verify all Data with Device Data Sheet and Product Availability with your local Xilinx Rep

Xilinx Software

	Feature	ISE *WebPACK*	ISE *BaseX*	ISE *Foundation*	ISE *Alliance*
	Platform	PC	PC	PC/UNIX	PC/UNIX
Device Support	Virtex Series	Up to 300K (Virtex E and Virtex II)	Up to 300K	ALL	ALL
	Spartan Series	Spartan-II Series	ALL	ALL	ALL
	XC4000 Series	No	4KE & Newer	4KE & Newer	4KE & Newer
	CoolRunner Series	ALL	ALL	ALL	ALL
	XC9500 Series	ALL	ALL	ALL	ALL
Design Planning	Modular Design	No	Sold as an Option	Sold as an Option	Sold as an Option
	Educational Services	Yes	Yes	Yes	Yes
	Design Services	Yes	Yes	Yes	Yes
	Support Services	Web Only	Yes	Yes	Yes
Design Entry	Schematic Editor (Gate & Block Level HDL)	Yes	Yes	PC Only	No
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	Yes	PC Only	No
	CORE Generator	No	Yes	Yes	Yes
	Xilinx System Generator for Simulink	No	Sold as an Option	Sold as an Option	Sold as an Option
Synthesis	FPGA Express	No	Yes	PC Only	No
	Xilinx Synthesis Technology (XST)	Yes	Yes	Yes	No
	Synplify/Pro Integration	No	Yes	PC Only	PC Only
	Leonardo Integration	Yes	Yes	Yes	Yes
	ABEL	CPLD	CPLD	CPLD (PC Only)	No
Implementation Tools	iMPACT	Yes	Yes	Yes	Yes
	FloorPlanner	Yes	Yes	Yes	Yes
	Xilinx Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	Timing Improvement Wizard	No	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	LMG SmartModels	Yes	Yes	Yes	Yes
Verification	HDL Bencher	Yes	Yes	PC Only	No
	Model*Sim* Xilinx Edition (XE) (See www.xilinx.com for more information)	Model*Sim* XE Starter Included (Model*Sim* XE Sold as an Option)	Model*Sim* XE Starter Included (Model*Sim* XE Sold as an Option)	Model*Sim* XE Starter Included (Model*Sim* XE Sold as an Option)	Model*Sim* XE Starter Included (Model*Sim* XE Sold as an Option)
	Static Timing Analyzer	Yes	Yes	Yes	Yes
	Chipscope ILA	No	Sold as an Option	Sold as an Option	Sold as an Option
	FPGA Editor with Probe	No	Yes	Yes	Yes
	ChipViewer	Yes	Yes	Yes	Yes
	XPower (Power Analysis)	Yes	Yes	Yes	Yes
	3rd Party Simulator Support	Yes	Yes	Yes	Yes
IP/CORE	For more information on the complete list of Xilinx IP products, visit the Xilinx IP Center at http://www.xilinx.com/ipcenter				

Xilinx Configuration Storage Solutions

System ACE	Memory Density	Number of Components	Min board space	Compression	FPGA Config. Mode	Multiple Designs	Software Storage	Removable	IRL Hooks	Max Config. Speed	Non-Volatile Media
System ACE CF	up to 8 Gbit	2	25 cm^2	No	JTAG	Unlimited	Yes	Yes	Yes	30 Mbit/sec	CompactFlash
System ACE MPM	16 Mbit 32 Mbit 64 Mbit	1	12.25 cm^2	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash Memory
System ACE SC	16 Mbit 32 Mbit 64 Mbit	3	Custom	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash memory

PROM

	Density	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage 2.5V	I/O Voltage 3.3V
In-System Programming (ISP) Configuration PROMs										
XC18V256	256Kb			Y	Y		Y	3.3V	Y	Y
XC18V512	512Kb			Y	Y		Y	3.3V	Y	Y
XC18V01	1Mb			Y	Y		Y	3.3V	Y	Y
XC18V02	2Mb					Y	Y	3.3V	Y	Y
XC18V04	4Mb					Y	Y	3.3V	Y	Y
One-Time Programmable (OTP) Configuration PROMs										
XC17V01	1.6Mb		Y	Y	Y			3.3V	Y	Y
XC17V02	2Mb				Y	Y	Y	3.3V	Y	Y
XC17V04	4Mb				Y	Y	Y	3.3V	Y	Y
XC17V08	8Mb					Y	Y	3.3V	Y	Y
XC17V16	16Mb					Y	Y	3.3V	Y	Y

	FPGA	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage 2.5V	I/O Voltage 3.3V
OTP Configuration PROMs for Spartan-II/IIE										
XC17S50A	XC2S50E	Y	Y	Y				3.3V	Y	Y
XC17S100A	XC2S100E	Y	Y	Y				3.3V	Y	Y
XC17S200A	XC2S150E	Y	Y				Y	3.3V	Y	Y
XC17S200A	XC2S200E	Y	Y				Y	3.3V	Y	Y
XC17S300A	XC2S300E						Y	3.3V	Y	Y
XC17S15A	XC2S15	Y	Y	Y				3.3V	Y	Y
XC17S30A	XC2S30	Y	Y	Y				3.3V	Y	Y
XC17S50A	XC2S50	Y	Y	Y				3.3V	Y	Y
XC17S100A	XC2S100	Y	Y	Y				3.3V	Y	Y
XC17S150A	XC2S150	Y	Y	Y				3.3V	Y	Y
XC17S200A	XC2S200	Y	Y				Y	3.3V	Y	Y

Xilinx Home Page
http://www.xilinx.com

Xilinx Online Support
http://www.xilinx.com/support/support.htm

Xilinx IP Center
http://www.xilinx.com/ipcenter/index.htm

Xilinx Education Center
http://www.xilinx.com/support/education-home.htm

Xilinx Tutorial Center
http://www.xilinx.com/support/techsup/tutorials/index.htm

Xilinx WebPACK
http://www.xilinx.com/sxpresso/webpack.htm

Xilinx IP Selection Guide

Function	Vendor Name	IP Type	Virtex-II	Virtex	Spartan-II	Spartan	Implementation Example: Occupancy	MHz	Device	Key Features	Application Examples
Communication & Networking											
BUFE-based Multiplexer Slice	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
3G FEC Package	Xilinx	LogiCORE	V-II	V						Viterbi Decoder, Turbo Codec, Convolutional Enc	3G Wireless Infrastructure
3GPP Compliant Turbo Convolutional Decoder	Xilinx	LogiCORE	V-II	V			80%	40	XC2V500	3GPP specs, 2 Mbps, BER=10-6 for 1.5dB SNR	3G Wireless Infrastructure
3GPP Compliant Turbo Convolutional Encoder	Xilinx	LogiCORE	V-II	V			65%	60	XC2V250	Compliant w/ 3GPP, puncturing	3G Wireless Infrastructure
3GPP Turbo Decoder	sysonchip	AllianceCORE	V-II	V			87%	66	XC2V500-5	3GPP/UMTS compliant, IMT-2000, 2Mbps data	Error correction, wireless
8b/10b Decoder	Xilinx	LogiCORE	V-II	V	S-II		1 BRAM	100	XC2V1000	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
8b/10b Encoder	Xilinx	LogiCORE	V-II	V	S-II		1 BRAM	100	XC2V1000	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
ADPCM 1024 Channel	Amphion	AllianceCORE	V-II	V						G.721, 723, 726, 726a, 727, 727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM 16 Channel	Amphion	AllianceCORE		V	S-II		89%	16	XCV150-6	G.721, 723, 726, 726a, 727, 727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM 256 Channel	Amphion	AllianceCORE	V-II	V						G.721, 723, 726, 726a, 727, 727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM 512 Channel	Amphion	AllianceCORE	V-II	V							
ADPCM 768 Channel	Amphion	AllianceCORE	V-II	V			89%	50	XC2V500-5	G.721, 723, 726, 726a, 727, 727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM Speech Codec, 32 Channel (DO-DI-ADPCM32)	Xilinx	LogiCORE	V-II	V	S-II		62%	25	XC2V500	G.726, G.727, 32 duplex channels	DECT, VOIP, Wireless local loop, DSLAM, PBX
ADPCM Speech Codec, 64 Channel (DO-DI-ADPCM64)	Xilinx	LogiCORE	V-II	V	S-II		61%	27	XC2V500	G.726, G.727, 64 duplex channels	DECT, VOIP, Wireless local loop, DSLAM, PBX
BOOST Lite Bluetooth Baseband Processor	NewLogic	AllianceCORE	V-II	V			73%	33	XC2V1000-4	Compliant to Bluetooth v1.1, BQB qualified software for L2CAP, LHP, HC1, voice support	Bluetooth applications
Cell Assembler (CC-201)	Paxonet	AllianceCORE		V		S	44%	60	XC4005XL-1	Octet wide operation, HEC compute, cell scrambling	ATM adapter cards, routers, switches
Cell Delineation (CC-200)	Paxonet	AllianceCORE		V		S	67%	40	XC4010XL-9	Octet wide operation, HEC verification, cell scrambling	ATM adapter cards, rtouters, switches
Convolutional Encoder	TILAB	AllianceCORE		V	S-II	S	2%	144	XCV50-6	code rate, gen. vectors, CMSTR length customizable	Error correction
Convolutional Encoder	Xilinx	LogiCORE	V-II	V	S-II		10%	26	XC2V40	k from 3 to 9, puncturing from 2/3 to 12/13	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave
CRC10 Generator and Verifier (CC-130)	Paxonet	AllianceCORE		V		S	22%	20	XCS30-4	Separate generator and verifier blocks, compatible with ITU-T I.363 for AAL3/AAL4	ATM, SONET, and Ethernet
CRC32 Generator and Verifier (CC-131)	Paxonet	AllianceCORE		V		S	43%	29	XCS30-4	Separate generator and verifier blocks, compat with ITU-T I.363 for AAL5	ATM, SONET, and Ethernet
DES	MemecCore	AllianceCORE		V	S-II	S	79%	25	XC520-4	NIST certified, supports EBC, CBC, CFB, and OFB	Secure communication, data storage
DES - Triple DES Cryptoprocessor	inSilicon	AllianceCORE		V	S-II	S	93%	48	XC2S150-6	Compliant with ANSI X9.52, 128-bit key or two independent 64-bit keys	Secure communication, data storage
DES Cryptoprocessor	inSilicon	AllianceCORE		V	S-II	S	20%	94	XC2S100-6	NIST certified, supports ECB, CBC, CFB, OFB	Secure communication, data storage
Distributed Sample Descrambler	TILAB	AllianceCORE		V	S-II	S	14%	74	XCV50-6	ITU-T I.432. Param data width, cell & header length	ATM PHY layer
Distributed Sample Descrambler	TILAB	AllianceCORE		V	S-II	S	14%	74	XCV50-6	ITU-T I.432. Param data width, cell & header length	ATM PHY layer
Distributed Sample Scrambler	TILAB	AllianceCORE		V	S-II	S	9%	104	XCV50-6	Compliant with ITU-T I.432 scrambler. Param data width, cell length, header length	ATM PHY layer
DVB Satellite Modulator Core	MemecCore	AllianceCORE		V	S-II	S	39%	45–70	XCV50-4	Conforms to ETSI EN 300 421 v1.1.2, selectable convolutional code rate	Digital broadcast, microwave transmitter
DVB-RCS Turbo Decoder	iCODING	AllianceCORE	V-II	V			54%	69	XC2V2000-5	DVB-RCS compliant, 9Mbps, data rate, switchable code rates and frame sizes	Error correction, wireless, DVB, Satellite data link
Fast Ethernet (10/100 Mbps) MAC Evaluation Board	Paxonet	AllianceCORE		V	S-II	S					
Fast Ethernet (10/100 Mbps) Media Access Controller Transmitter and Receiver Cores	Paxonet	AllianceCORE		V	S-II	S	90%	50	XCV150-4	IEEE 802.3 compliant RMON, MIBs stats, MII	Ethernet switched, hub, NICS
Flexbus 4 Interface Core, 16-Channel (DO-DI-FLX4C16)	Xilinx	LogiCORE	V-II				31%	200	XC2V3000 FG676-5		Line card: terabit routers & optical switches
Flexbus 4 Interface Core, 4-Channel (DO-DI-FLX4C4)	Xilinx	LogiCORE	V-II				27%	200	XC2V1000 FG456-5		Line card: terabit routers & optical switches
Flexbus 4 Interface Core, 1-Channel (DO-DI-FLX4C1)	Xilinx	LogiCORE	V-II				12%	200	XC2V1000 FG456-5		Line card: terabit routers & optical switches
G.711 PCM Codec	Xilinx	LogiCORE		V	S-II		12%	44	XCV50	µ-Law, ITU G.711, EBI for A-Law	Digital telephony, DECT, T1 & E1 Links
G.711 PCM Compressor	Xilinx	LogiCORE		V	S-II		7%	44	XCV50	µ-Law, ITU G.711, EBI for A-Law	Digital telephony, DECT, T1 & E1 Links
G.711 PCM Expander	Xilinx	LogiCORE		V	S-II		6%	57	XCV50	Digital telephony, DECT, T1 & E1	Digital telephony, DECT, T1 & E1 Links
HDLC Controller Core, 32 Channels	Xilinx	LogiCORE	V-II	V	S-II		34%	81	XC2V250	32 full duplex, CRC-16/32, 8/16-bit address insertion/deletion	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
HDLC Controller Core, Single Channel	Xilinx	LogiCORE	V-II	V	S-II		15%	115	XC2V250	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detect	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
IMA-32 Inverse Multiplexer for ATM	Mindspeed	AllianceCORE		V			78%	50	XCV400E-6	Compliant ATM Forum IMA, prog. groups & links, SW driver available	Network adapters, routers, multiplexers
IMA-8 Inverse Multiplexer for ATM	Mindspeed	AllianceCORE			S-II		100%	50	XC2S150-5	Compliant with ATM Forum IMA, prog. groups & links, SW driver	Network adapters, routers, DSLAMs
Interleaver Deinterleaver	TILAB	AllianceCORE		V	S-II	S	21%	73	XVC50-6	Block & convolutional support. param features. 3GPP, UMTS, GSM, DVB compliant	Channel coding in telecom/wireless, broadcast
Interleaver/De-interleaver	Xilinx	LogiCORE	V-II	V	S-II		30%	187	XC2V40	Convolutional, width up to 256 bits, 256 branches	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
IPlogiCAM Internet Protocol Content Addressable Memory	TILAB	AllianceCORE		V	S-II	S	9%	49	XCV50-6	Hardware control blk works with s/w CAM	IP routers
MT1F T1 Framer	Virtual	AllianceCORE					*	*	*	D4, ESF, SLC-96 formats. For XC4000.	DSI trunk, PBX I/F
Noisy Transmission Channel Model	TILAB	AllianceCORE		V	S-II	S	22%	100	XCV50-6	Programmable noise generation profile	Noise emulation in transmission channel
PARSER: Bit Stream Analyzer and Data Extractor	TILAB	AllianceCORE		V	S-II	S	32%	67	XCV50-6	Data syntax analysis of IP, MPEG, ATM	ATM, IP, MPEG
PE-MACMII Dual Speed 10/100 Mbps Ethernet MAC	Alcatel	AllianceCORE	V-II	V	S-II		33%	60	XC2V500-4	802.3 compliant, Supports single & multimode fiber optic devices, M11 interfaces, RMON and Etherstate statistics	Networking, Broadband, SOHO, Home networking, storage, routers, switches, printers, NIC
POS-PHY Level 3 Link Layer Interface Core, 48 Channel (DO-DI-POSL3LINK48A)	Xilinx	LogiCORE	V-II				33%	104	XC2V6000 FF1152-4		
POS-PHY L3 Link Layer Interface, 16-Ch (DO-DI-POSL3LINK16)	Xilinx	LogiCORE	V-II				40%	104	XC2V1000 FG456-4		Line card: terabit routers & optical switches
POS-PHY L3 Link Layer Interface, 4-Ch (DO-DI-POSL3LINK4)	Xilinx	LogiCORE	V-II	VE			15%	104	XC2V1000 FG456-4		Line card: terabit routers & optical switches
POS-PHY L3 Link Layer Interface, 2-Ch (DO-DI-POSL3LINK2)	Xilinx	LogiCORE	V-II	VE			55%	104	XCV50E-8		Line card: terabit routers & optical switches
POS-PHY L3 Link Layer Interface, Single Channel	Xilinx	LogiCORE	V-II	VE			6%	104	XC2V1000 FG456-4		
POS-PHY L3 Physical Layer Interface (DO-DI-POSL3PHY)	Xilinx	LogiCORE		VE			52%	104	XCV50E-8		Line card: terabit routers & optical switches
POS-PHY L4 Multi-Channel Interface (DO-DI-POSL4MC)	Xilinx	LogiCORE	V-II				29%	104	XC2V3000 FG676-5		
PPP8 HDLC Core CC318f	Paxonet	AllianceCORE		V	S-II		76%	80	XC2S150-6	RFC1619 (IP&IPX) POS, 16/32 bit FCS generation and verification, stats	Bridges, switches, WAN links
Reed-Solomon Decoder	MemecCore	AllianceCORE		V	S-II	S	83%	73	XCV50-6	Customizable, >580 Mbps	Error correction
Reed-Solomon Decoder	Amphion	AllianceCORE		V	S-II	S	51%	50	XCV100-4	Supports ETSI 300 421, 300 429, >300 Mbps	Error correction
Reed-Solomon Encoder	MemecCore	AllianceCORE		V	S-II	S	13%	113	XCV50-6	Customizable, > 900 Mbps	Error Correction
Reed-Solomon Encoder	Amphion	AllianceCORE		V	S-II	S	11%	82	XCV50-4	Supports ETSI 300 421, 300 429	Error correction
Reed-Solomon Decoder	Xilinx	LogiCORE	V-II	V	S-II	S	40%	98	XC2V250	Std or custom coding, 3-12 bit symbol width, up to 4095 symbols	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV

Function	Vendor Name	IP Type	Virtex-II	Virtex	Spartan-II	Spartan	Implementation Example: Occupancy	Implementation Example: MHz	Implementation Example: Device	Key Features	Application Examples
Microprocessors, Controllers & Peripherals (continued)											
OPB Timer/Counter	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
OPB UART (16450, 16550)	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Interfaces through OPB to MicroBlaze	Processor applications
OPB UART Lite	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
OPB WDT	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
PF3100 PC/104-Plus Reconfigurable Module	Derivation	AllianceCORE	V-II				N/A	N/A	XC2V1000 FG256	PC/104 & PC/104+ devlopment board	Internet appliance, industrial control
R8051 RISC MicroController	CAST	AllianceCORE		V	S-II		76%	34	XC2S150-6	12X faster, SRF I/F	Embedded systems
R80515 High-speed 8-bit RISC Microcontroller	CAST	AllianceCORE		V			56%	42	XCV200E-8	RISC implementation, 8 bit ALU, 8 bit control, 32 bit I/O, 16 bit timer/counters, SFR I/F, ext. memory I/F	High speed embedded systems, audio, video
SPI	Xilinx	LogiCORE	V-II	V	S-II				Virtex-II	Interfaces through OPB to MicroBlaze	Networking, communications, processor apps
Synchronous DRAM Controller	NMI	AllianceCORE		V	S-II		5%	137	XCV50-6	SDRAM refresh, customizable	Embedded systems using SDRAMs
uEBX Reference and Development Platform	NMI	AllianceCORE		V			NA	NA	NA	Interfaces NMI's MicroEngines to EBX bus	PC104 applications
uPCI Reference and Development Platform	NMI	AllianceCORE		V			NA	NA	NA	Interfaces NMI's MicroEngines to PCI bus	PCI ethernet/graphics applications
V8-uRISC 8-bit RISC Microprocessor	VAutomation	AllianceCORE					*	*	*	8-bit processor, 8 bit ALU, 16 bit stack pointer, 33 opcodes, 4 addr Modes, 2 user opcodes	Embedded systems, 8-bit processing apps.
XF8250 UART	MemecCore	AllianceCORE				S	59%	10	XCS10-4	DC to 625K baud	Serial communications
XF8255 Programmable Peripheral Interface	MemecCore	AllianceCORE				S	64%	8	XCS05-4	Bit set/reset support	Embedded systems
XF8256 Multifunction Microprocessor Support Controller	MemecCore	AllianceCORE				S	89%	10	XCS20-4	Baud rate generator for 13 common baud rates, parallel I/O ports, prog. timer/counters	Communication, embedded systems
XF8279 Programmable Keyboard Display Interface	MemecCore	AllianceCORE				S	46%	8	XCS20-4	8 char keyboard FIFO, 2-key lockout, n-key rollover, 4-16 char display	Embedded systems interface
XF-TWSI Two-Wire Serial Interface Master-Only	MemecCore	AllianceCORE		V	S-II	S					
XF-TWSI-MS Two-Wire Serial Interface Master-Slave	MemecCore	AllianceCORE		V	S-II	S	24%	59	XCV50-4	I2C-like, multi master fast/std. modes	Embedded systems
XF-UART Asynchronous Communications Core	MemecCore	AllianceCORE	V-II	V		S	15%	50	XCS20-4	UART and baud rate generator	Serial data communication
Standard Bus Interfaces											
Arbiter	TILAB	AllianceCORE		V	S-II	S		33	XCV50-6	2 priority classes - strong/weak, access counters	General purpose bus arbitration
CAN Bus Interface R3.0	Sci-worx	AllianceCORE		V	S-II	S	*	*	*	Supports CAN 2.0A, 2.0B, error handling, stuff bit generation, SRC, individual acceptance filtering	Automotive, network, home automation
EP100 PowerPC Bus Slave	Eureka	AllianceCORE		V	S-II						
EP201 PowerPC Bus Master	Eureka	AllianceCORE		V	S-II						
PCI-X 64/66 Interface for Virtex-E (DO-DI-PCIX64-VE)	Xilinx	LogiCORE		VE			30%	66	XCV300E-8	PCI-X 1.0 comp, 64/32-bit, 66 MHz PCI-X initiator and target IF, PCI 2.2 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	Comm systems, SAN, clustered servers, Ultra 3 SCSI/Fibre Ch RAID, multi-port Gb
PCI-X 64/100 Interface for Virtex-II (DO-DI-PCIX64-VE)	Xilinx	LogiCORE	V-II				30%	100	XC2V1000 FG456-5	PCI-X 1.0 comp, 64/32-bit, 66 MHz PCI-X initiator and target IF, PCI 2.2 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	Server, Embedded, gb ethernet, U320 SCSI, Fibre Ch, RAID cntl, graphics
PCI32 Single-Use License for Spartan (DO-DI-PCI32-SP)	Xilinx	LogiCORE			S-II	S	12%	66	XC2S200 PQ208-6	v2.2 comp, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI32 Virtex Interface Design Kit (DO-DI-PCI32-DKT)	Xilinx	LogiCORE	V-II	V	S-II	S	6%	66	XC2V1000 FG456-5	Includes PCI32 board, drive development kit, and customer education 3-day training class	
PCI32 Virtex Interface, IP Only (DO-DI-PCI32-IP)	Xilinx	LogiCORE	V-II	V	S-II	S	6%	66	XC2V1000 FG456-5	v2.2 comp, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI64 & PCI32, IP Only (DO-DI-PCI-AL)	Xilinx	LogiCORE	V-II	V	S-II	S	6 - 7%	66	XC2V1000 FG456-5	v2.2 comp, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards, CPCI, Embedded, hiperf video, gb ethernet
PCI64 Virtex Interface Design Kit (DO-DI-PCI64-DKT)	Xilinx	LogiCORE	V-II	V	S-II	S	7%	66	XC2V1000 FG456-5	v2.2 comp, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards, CPCI, Embedded, hiperf video, gb ethernet
PCI64 Virtex Interface, IP Only (DO-DI-PCI64-IP)	Xilinx	LogiCORE	V-II	V	S-II		7%	66	XC2V1000 FG456-5	v2.2 comp, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards, CPCI, Embedded, hiperf video, gb ethernet
RapidIO 8-bit port LP-LVDS Phy Layer (DO-DI-RIO8-PHY)	Xilinx	LogiCORE	V-II				24%	250	XC2V1000 FG456-5	RapidIO Interconnect v1.1 compliant, verified with Motorola's RapidIO bus functional model v1.4	Routers, switches, backplane, control plane, data path, embedded sys, high speed interface to memory and encryption engines, high end video
USB 1.1 Device Controller	MemecCore	AllianceCORE	V-II	V	S-II		21%	12	XC2V1000-5	Compliant with USB1.1 spec., Supports VCI bus, Performs CRC, Supports 1.5 Mbps & 12 Mbps	Scanners, Printers, Handhelds, Mass Storage
Video & Image Processing											
1-D Discrete Cosine Transform	Xilinx	LogiCORE	V-II	V	S-II					8-24 bits for coeff & input, 8-64 pts	
2-D DCT/IDCT Forward/Inverse Discrete Cosine Transform	Xilinx	LogiCORE	V-II	V	S-II						image, video phone, color laser printers
DCT/IDCT Forward/Inverse Discrete Cosine Transform	inSilicon	AllianceCORE		V	S-II	S	86%	29	XC2S100-6	DCT & IDCT, one clock cycle per sample. 38 MHz when operated as DCT	JPEG, MPEG, H261 designs
FASTJPEG_BW DECODER	BARCO-SILEX	AllianceCORE	V-II	V			67%	73	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, Gray-Scale	Video editing, digital camera, scanners
FASTJPEG_C DECODER	BARCO-SILEX	AllianceCORE	V-II	V			78%	56	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, color, multi-scan, Gray-Scale	Video editing, digital camera, scanners
FIDCT Forward/Inverse Discrete Cosine Transform	TILAB	AllianceCORE		V	S-II		77%	78	XCV200-6	DCT for 8X8, 16X16, IDCT IEEE1180-1990 compliant	JPEG, MPEG, H.26X
JPEG CODEC	inSilicon	AllianceCORE		V	S-II	S	75%	20	XCV400E-8	Conforms to ISO/IEC Baseline 10918-1, 4 quantization tables, 4 Huffman tables. Stallable	Video editing, digital camera, scanners
logiCVC - Compact Video Controller	xylon	AllianceCORE		V	S-II		35%	88	XC2V250-4	Single & double panel, LCD/CRT support, 4 gray, 256 colors	Video phone, Set-top box, PDA display
Longitudinal Time Code Generator	DELTATEC	AllianceCORE		V			12%	85	XCV100-4	SMTPE/EBU compliant, PAL/NTSC, lock-on external video reference	Audio/Video recording and editing equipment
RGB2YCrCb Color Space Converter	Perigee	AllianceCORE		V	S-II	S	22%	202	XCV100E-8	One clock cycle throughput	HDTV, real time TV output modulation
RGB2YCrCb Color Space Converter Core	Xilinx	LogiCORE		V	S-II		49%	65	XC2S30		TV, HDTV, color imaging, color video
RGB2YUV Color Space Converter Core	Xilinx	LogiCORE		V	S-II		53%	80	XC2S30		TV, HDTV, color imaging, color video
YCrCb2RGB Color Space Converter	Perigee	AllianceCORE		V	S-II	S	16%		XCV100E-8	One clock cycle throughput	HDTV, real time video
YCrCb2RGB Color Space Converter	Xilinx	LogiCORE		V	S-II		15%	70	XC2S100		TV, HDTV, color imaging, color video
YUV2RGB Color Space Converter	Xilinx	LogiCORE		V	S-II		12%	65	XC2S100		TV, HDTV, color imaging, color video
Basic Elements											
BUFE-based Multiplexer Slice	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
BUFT-based Multiplexer Slice	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Binary Counter	Xilinx	LogiCORE	V-II	V	S-II					2-256 bits output width	
Binary Decoder	Xilinx	LogiCORE	V-II	V	S-II					2-256 bits output width	
Bit Bus Gate	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Bit Gate	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Bit Multiplexer	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Bus Gate	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Bus Multiplexer	Xilinx	LogiCORE	V-II	V	S-II					IO widths up to 256 bits	
Comparator	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
FD-based Parallel Register	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
FD-based Shift Register	Xilinx	LogiCORE	V-II	V	S-II					1-64 bits wide	
Four-Input MUX	Xilinx	LogiCORE				S					
LD-based Parallel Latch	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide	
Parallel-to-Serial Converter	Xilinx	LogiCORE				S					
RAM-based Shift Register	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits wide, 1024 words deep	
Register	Xilinx	LogiCORE				S					
Three-Input MUX	Xilinx	LogiCORE				S					
Two-Input MUX	Xilinx	LogiCORE				S					

Function	Vendor Name	IP Type	Virtex-II	Virtex	Spartan-II	Spartan	Occupancy	MHz	Device	Key Features	Application Examples
Communication & Networking (continued)											
Reed-Solomon Decoder	TILAB	AllianceCORE	V-II	V			56%	61	XC2V1000-5	parameterizable, RTL available	Error correction, wireless, DSL
Reed-Solomon Encoder	Xilinx	LogiCORE	V-II	V	S-II	S	42%	180	XC2V40	Std or cust coding, 3-12 bit width, up to 4095 symbols with 256 check symb	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
SDLC Controller	CAST	AllianceCORE	V-II	V			38%	158	XC2V100-5	Like Intel 8XC152 Global Serial Channel, Serial Comm., HDLC apps, telecom	Embedded systems, professional audio, video
Single-Channel XF-HDLC Controller	MemecCore	AllianceCORE		V	S-II		95%	77	XC2S15-5	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detection	X.25, Frame Relay, B/D-Channel
SPEEDROUTER Network Processor	IP	AllianceCORE	V-II	V			64%	80, 2.5Gbps	XC2V1500-5	Solution requires SPEEDAnalyzer ASIC, 2.5 Gbps fdx wire speed; net processor (NPV)	Networking, edge and access, Switches and routers
T1 Deframer	Xilinx	LogiCORE		V	S-II		15%	54	XC2S150		ISDN PRA links, mux equip, satellite com, digital PABX, high-speed computer links
T1/E1 Framer	Xilinx	LogiCORE		V	S-II		7%	72	XC2S150		ISDN PRA links, mux equip, satellite com, digital PABX, high-speed computer links
Turbo Decoder - 3GPP	SysOnChip	AllianceCORE	V-II	V			88%	65	XC2V2000-5	3GPP/UMTS compliant, 2Mbps data rate	Error correction, wireless
Turbo Encoder	TILAB	AllianceCORE	V-II	V	S-II		48%	120	XC2V80-5	3GPP/UMTS compliant, upto 4 interleaver laws	Error correction, wireless
TURBO_DEC Turbo Decoder	TILAB	AllianceCORE	V-II	V			99%	65	XC2V2000-5	3GPP/UMTS compliant, >2Mbps data rate	Error correction, wireless
UTOPIA Level-2 PHY Side RX Interface	TILAB	AllianceCORE		V	S-II	S	8%	53	XCV50-6	Protocol conversion from Pb (RACE BLNT) to UTOPIA L2, 8/16 bit operation	ATM PHY layer
UTOPIA Level-2 PHY Side TX Interface	TILAB	AllianceCORE		V	S-II	S	10%	61	XCV50-6	Protocol conversion from UTOPIA L2 Pb (RACE BLNT), 8/16 bit operation	ATM PHY layer
UTOPIA Level-3 ATM Receiver	inSilicon	AllianceCORE		V	S-II	S	5%	164	XCV100E-8	Supports ATM Forum UTOPIA Level-3. Configurable cell format, data width	High capacity ATM switches
UTOPIA Level-3 ATM Transmitter	inSilicon	AllianceCORE		V	S-II	S	6%	150	XCV100E-8	Supports ATM Forum UTOPIA Level-3. Configurable cell format, data width	High capacity ATM switches
UTOPIA Level-3 PHY Receiver	inSilicon	AllianceCORE		V	S-II	S	21%	104	XCV100E-8	Supports ATM Forum UTOPIA Level-3. Configurable cell format, data width, configurable FIFO size	High capacity ATM switches
UTOPIA Level-3 PHY Transmitter	inSilicon	AllianceCORE		V	S-II	S	22%	100	XCV100E-8	Supports ATM Forum UTOPIA Level-3. Configurable cell format, data width, configurable FIFO size	High capacity ATM switches
UTOPIA Master (CC140f)	Paxonet	AllianceCORE		V	S-II	S	*	*	*	SPHY, MPHY, HEC processing, round robin polling, ind. transmitter receiver	ATM PHY layer
UTOPIA Slave (CC141)	Paxonet	AllianceCORE				S	*	*	*	Cell handshake in SPHY mode, 8/16 bit operation, internal FIFO, detects runt cells	ATM PHY layer
UTOPIA Slave (CC143S)	Paxonet	AllianceCORE		V	S-II	S	26%	79	XCV50-4	Cell handshake in SPHY mode, 8/16 bit operation, internal FIFO, detects runt cells	ATM PHY layer
Viterbi Decoder	TILAB	AllianceCORE		V	S-II	S	65%	56	XCV50-6	Radix-2/radix4 architectures, BER, depuncturing. Code rate, contraint length parameterizable	Data transmission, wireless
Viterbi Decoder	Xilinx	LogiCORE	V-II	V	S-II		80%	100	XC2V250	Puncturing, serial & parallel architecture,	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave
Viterbi Decoder, IEEE 802-compatible	Xilinx	LogiCORE	V-II	V	S-II		70%	147	XC2V250	Constraint length(k)=7, G0=171, G1=133	LMMDS, broadcast equip, wireless LAN, cable modem, xDSL, sat com, uwave nets
Digital Signal Processing											
1024-Point Complex FFT IFFT	Xilinx	LogiCORE		V							
1024-Point Complex FFT IFFT for Virtex-II	Xilinx	LogiCORE	V-II				62%	100, 41us	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
16-Point Complex FFT / IFFT	Xilinx	LogiCORE	V-II	V							
16-Point Complex FFT IFFT for Virtex-II	Xilinx	LogiCORE	V-II				37%	130, 123ns	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
256-Point Complex FFT / IFFT	Xilinx	LogiCORE	V-II	V							
256-Point Complex FFT IFFT for Virtex-II	Xilinx	LogiCORE	V-II				54%	100, 7.7us	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
32 Point Complex FFT/IFFT	Xilinx	LogiCORE	V-II	V	S-II						
64-Point Complex FFT IFFT	Xilinx	LogiCORE		V							
64-Point Complex FFT IFFT for Virtex-II	Xilinx	LogiCORE	V-II				38%	100, 1.9us	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
Bit Correlator	Xilinx	LogiCORE	V-II	V	S-II					4096 taps, serial/parallel input, 4096 bits width	
Cascaded Integrator Comb (CIC)	Xilinx	LogiCORE	V-II	V	S-II					32 bits data width, rate change from 8 to 16384	
Comb Filter	Xilinx	LogiCORE				S					
Direct Digital Synthesizer	Xilinx	LogiCORE	V-II	V	S-II					8-65K samples, 32-bits output precision, phase dithering/offset	
Distributed Arithmetic FIR Filter	Xilinx	LogiCORE	V-II	V	S-II					32-bit input/coeff width, 1024 taps, 1-8 chan, polyphase, online coeff reload	
Dual-Channel Numerically Controlled Oscillator	Xilinx	LogiCORE		V	S-II	S					
GVA-200A DSP Hardware Accelerator	GV	AllianceCORE				S	NA	NA	*	Ext SRAM I/F	DSP prototyping
GVA-220 DSP Hardware Accelerator	GV	AllianceCORE				S	NA	NA	*	Ext SRAM I/F, 2 FPGA's	DSP prototyping
GVA-250 Virtex DSP Hardware Accelerator	GV	AllianceCORE		V			NA	NA	*	2 FPGAs, 2 SRAMs	DSP prototyping
GVA-270 Virtex-E DSP Hardware Accelerator	GV	AllianceCORE		V			NA	NA	*	Virtex-E support, 2 FPGAs, ZBT SRAMs	DSP prototyping
GVA-290 Virtex-E DSP Hardware Accelerator	GV	AllianceCORE		V			NA	NA	*	2 Virtex-E, Spartan-II FPGAs, 1 CPLD, Matlab I/F	DSP prototyping
GVA-300 Virtex-II DSP Hardware Accelerator	GV	AllianceCORE	V-II				NA	NA	*	2 Virtex-II, Spartan-II FPGAs, 1 CPLD, Matlab I/F	DSP prototyping
LFSR, Linear Feedback Shift Register	Xilinx	LogiCORE	V-II	V	S-II					168 input widths, SRL16/register implementation	
Nonsymmetric 16-Deep Time-Skew Buffer	Xilinx	LogiCORE				S					
Nonsymmetric 32-Deep Time-Skew Buffer	Xilinx	LogiCORE				S					
Numerically Controlled Oscillator	Xilinx	LogiCORE		V	S-II	S					
Parallel Distributed Arithmetic FIR Filter	Xilinx	LogiCORE				S					
Serial Distributed Arithmetic FIR Filter	Xilinx	LogiCORE				S					
Symmetric 16 Deep Time-Skew Buffer	Xilinx	LogiCORE				S					
Math Functions											
1s and 2s Complement	Xilinx	LogiCORE				S					
Accumulator	Xilinx	LogiCORE	V-II	V	S-II					1-256s bit wide	
Adder Subtracter	Xilinx	LogiCORE	V-II	V	S-II					1-256s bit wide	
Constant Coefficient Multiplier	Xilinx	LogiCORE				S					
Constant Coefficient Multiplier - Pipelined	Xilinx	LogiCORE				S					
DFP2INT Floating Point to Integer Converter	Digital	AllianceCORE	V-II	V	S-II		39%	66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
DFPADD Floating Point Adder	Digital	AllianceCORE	V-II	V	S-II		39%	66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
DFPCOMP Floating Point Comparator	Digital	AllianceCORE	V-II	V	S-II		16%	91	XC2V80-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps.
DFPDIV Floating Point Divider	Digital	AllianceCORE	V-II	V	S-II		99%	53	XC2V250-5	Full IEEE-754 compliance, 15 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
DFPMUL Floating Point Multiplier	Digital	AllianceCORE	V-II	V	S-II		44%	74	XC2V250-5	Full IEEE-754 compliance, 7 pipelines, 32x32 mult, Single precision real format support	DSP, Math, Arithmetic apps.
DFPSQRT Floating Point Square Root	Digital	AllianceCORE	V-II	V	S-II		39%	66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
DINT2FP Integer to Floating Point Converter	Digital	AllianceCORE	V-II	V	S-II		37%	73	XC2V250-5	Full IEEE-754 compliance, double word input, 2 pipelines, Single precision real output	DSP, Math, Arithmetic apps
Dynamic Constant Coefficient Multiplier	Xilinx	LogiCORE		V							
Integrator	Xilinx	LogiCORE				S					
Multiply Accumulator (MAC)	Xilinx	LogiCORE	V-II	V	S-II					Input width up to 32 bits, 65-bit accumulator, truncation rounding	
Multiply Generator	Xilinx	LogiCORE	V-II	V	S-II					64-bit input data width, constant, reloadable or variable inputs, parallel/sequential implementation	
Parallel Multipliers Area Optimized	Xilinx	LogiCORE				S					
Pipelined Divider	Xilinx	LogiCORE	V-II	V	S-II	S				32-bit input data width, multiple clock per output	
Registered Adder	Xilinx	LogiCORE				S					
Registered Loadable Adder	Xilinx	LogiCORE				S					

Function	Vendor Name	IP Type	Virtex-II	Virtex	Spartan-II	Spartan	Implementation Example: Occupancy	MHz	Device	Key Features	Application Examples
Math Functions (continued)											
Registered Loadable Subtracter	Xilinx	LogiCORE				S					
Registered Scaled Adder	Xilinx	LogiCORE				S					
Registered Serial Adder	Xilinx	LogiCORE				S					
Registered Subtracter	Xilinx	LogiCORE				S					
Scaled-by-One-Half Accumulator	Xilinx	LogiCORE				S					
Sine Cosine Look Up Table	Xilinx	LogiCORE	V-II	V	S-II	S				3-10 bit in, 4-32 bit out, distributed/block ROM	
Square Root	Xilinx	LogiCORE				S					
Twos Complementer	Xilinx	LogiCORE	V-II	V	S-II					Input width up to 256 bits	
Variable Parallel Virtex Multiplier	Xilinx	LogiCORE		V	S-II						
Memories & Storage Elements											
Asynchronous FIFO	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits, 15-65535 words, DRAM or BRAM, independent I/O clock domains	
Content Addressable Memory (CAM)	Xilinx	LogiCORE	V-II	V	S-II					1-512 bits, 2-10K words, SRL16	
Distributed Memory	Xilinx	LogiCORE	V-II	V	S-II					1-1024 bit, 16-65536 word, RAM/ROM/SRL16, opt output regs and pipelining	
Dual-Port Block Memory	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits, 2-13K words	
Pipelined Delay Element	Xilinx	LogiCORE				S					
Registered ROM	Xilinx	LogiCORE				S					
Registered Single Port RAM	Xilinx	LogiCORE				S					
Single-Port Block Memory	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits, 2-128K words	
Synchronous FIFO	Xilinx	LogiCORE				S					
Synchronous FIFO	Xilinx	LogiCORE	V-II	V	S-II					1-256 bits, 16-256 words, distributed/block RAM	
Microprocessors, Controllers & Peripherals											
10/100 Ethernet MAC	Xilinx	LogiCORE	V-II	V	S-II				Virtex-II	Interfaces through OPB to MicroBlaze	Networking, communications, processor apps
16-Word-Deep Registered Look Up Table	Xilinx	LogiCORE				S					
200 MHz SDRAM Controller Core	Rapid	AllianceCORE		V	S-II						
ARC 32-bit Configurable RISC Processor	ARC	AllianceCORE		V	S-II		89%	37	XC2S150-6	4 stage pipeline, 16 single cycle instructions10, 3 interrupt exception levels, 24 bit stack pointer	32 bit processing, DSP
AX1610 16-bit RISC Processor	Loarant	AllianceCORE	V-II	V	S-II		12%	91	XC2V500-5	44 opcode, 64-K word data, program, Harvard arch.	Control functions, State mach, Coprocessor
C16450 UART	CAST	AllianceCORE		V	S-II		29%	60	XC2S50-6	Independently controlled transmit, receive and data interrupts. 16X clock	Serial data applications, modems
C16550 UART with FIFOs	CAST	AllianceCORE		V	S-II			51	XC2S50-6	Prog. Data width, parity, stop bits. 16X internal clock, FIFO mode, false start bit detection	Serial data applications, modems
C165X MicroController	CAST	AllianceCORE	V-II	V	S-II		60%	134	XC2V80-5	Microchip 16C5X PIC like	Embedded systems, telecom
C2901 Microprocessor Slice	CAST	AllianceCORE		V	S-II		19%	36	XC2S50-6	Eight function ALU, 4 status flags- Carry, Overflow, Zero and Negative	Simple microcontroller applications
C2910a Microprogram Controller	CAST	AllianceCORE		V	S-II	S	13%	63	XCV50-6	Based on AMD 2910a	High-speed bit slice design
C68000 Microprocessor	CAST	AllianceCORE	V-II	V			90%	32	XC2V500-5	MC68000 Compatible	Embedded systems, pro audio, video
C8051 MicroController	CAST	AllianceCORE		V	S-II		52%	68	XCV200E-8	80C31 instruction set, 8 bit ALU, 8 bit control, 32 bit I/O ports, two 16 bit timer/counters, SFR I/F	Embedded systems, telecom
C8250 UART	CAST	AllianceCORE		V	S-II			98	XCV50E-8	UART & Baud rate generator, 16X clock generator, loopback & echo modes	Serial data applications, modems
C8254 Programmable Interval Timer/Counter	CAST	AllianceCORE		V	S-II		38%	66	XCV100E-8	Status feedback, counter latch, square wave mode, binary/BCD count, LSB/MSB R/W	Event counterr, baud rate generator
C8255A Peripheral Interface	CAST	AllianceCORE		V	S-II		10%	227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable IO lines, 8-bit bidi data bus	Processor I/O interface
C8259A Programmable Interrupt Controller	CAST	AllianceCORE		V	S-II		28%	47	XC2S50-6	8 vectored priority interrupts, all 8259/A modes programmable- e.g., special mask, buffer	Real-time interrupt based uP designs
Compact UART	CAST	AllianceCORE		V	S-II	S	7%	142	XCV50-6	1 start bit, 1 stop bit. Polling and interrupt modes	Serial data applications, modems
Compact Version of D80530 Microcontroller	CAST	AllianceCORE		V	S-II		88%	51	XC2S150-6	32 bit I/O, 3 counters, interrupt controller, SFR interface, dual data pointer	Low cost embedded systems, telecom
CPU + FPGA (Virtex/Spartan-II) MicroEngine Cards	NMI	AllianceCORE		V	S-II		NA	NA	NA	Hitachi SH-4 CPU	Embedded systems
CPU + FPGA (Virtex-II) MicroEngine Cards	NMI	AllianceCORE	V-II				NA	NA	NA	Hitachi SH-3 CPU	Embedded systems
CZ80CPU Microprocessor	CAST	AllianceCORE	V-II	V	S-II		55%	72	XC2V500-5	Zilog Z80 compatible, 8-bit processor	Embedded systems, Communications
D80530 8-bit Microcontroller	CAST	AllianceCORE		V			81%	66	XCV200E-8	32 bit I/O, 3 counters, 27-bit watchdog timer, 3-priority interrupt controller, SFR interface	Embedded systems, telecom
DDR SDRAM Controller Core	MemecCore	AllianceCORE	V-II	V	S-II		7%	133	XC2V1000-4	DDR SDRAM burst length support for 2,4, 8 per access, supports data 16, 32, 64, 72	Digital video, embedded computing , networking
DFPIC125X Fast RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		49%	126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
DFPIC1655X Fast RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		79%	140	XC2V80-5	S/W compatible with PIC16C55X, 14-bit instruction set, 35 instructions	Embedded systems, telecom, audio and video
DFPIC165X Fast RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		49%	126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
DI2CM I2C Bus Controller Master	Digital	AllianceCORE	V-II	V	S-II	S	58%	143	XC2V50-5	I2C-like, multi master, fast/std. modes	Embedded systems
DI2CM I2C Bus Controller Slave	Digital	AllianceCORE	V-II	V	S-II	S	28%	157	XC2V50-5	I2C-like, Slave	Embedded
DI2CSB I2C Bus Controller Slave Base	Digital	AllianceCORE	V-II	V	S-II	S	15%	187	XC2V50-5	I2C-like, Slave	Embedded Systems
DR8051 RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		68%	73	XC2V250-5	80C31 instruction set, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
DR8051BASE RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		46%	80–90	XC2V250-5	80C31 instruction set, high speed multiplier, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
DR8052EX RISC MicroController	Digital	AllianceCORE	V-II	V	S-II		99%	71	XC2V250-5	80C31 instruction set, high speed mult/div ,RISC 6.7X faster than standard 8051	Embedded systems, telecom, video
e8254 Programmable Interval Timer/Counter	einfochips	AllianceCORE	V-II	V	S-II		1%	175	XC2V1000-5	Three 8-bit parallel ports, 24 programmable IO lines, 8-bit bidi data bus	Processor, I/O interface
e8255 Peripheral Interface	einfochips	AllianceCORE	V-II	V	S-II		1%	175	XC2V1000-5	Three 8-bit parallel ports, 24 programmable IO lines, 8-bit bidi data bus	Processor, I/O interface
EP520 SDRAM Controller	Eureka	AllianceCORE		V	S-II						
Flip805x-PR Core	Dolphin	AllianceCORE		V	S-II	S	68%	20	XC2S150-6	Avg 12X faster and code compatible v. legacy 8051, verification bus monitor, SFR interface	Telecom, industrial, high speed control
Flip805x-PS Microprocessor	Dolphin	AllianceCORE	V-II	V	S-II		39%	38	XC2V1000-5	Avg 8X faster & code compatible v. legacy 8051, verification bus monitor, SFR IF, DSP focused	DSP, Telecom, industrial, high speed control
IIC	Xilinx	LogiCORE	V-II	V	S-II				Virtex-II	Interfaces through OPB to MicroBlaze	Networking, communications, processor apps
IntelliCore Prototyping System	VAutomation	AllianceCORE					*	*	*	USB, 1394, 1284, RS-232, IVDA I/F	Prototyping
LavaCORE Configurable Java Processor Core	Derivation	AllianceCORE	V-II	V			38%	20	XC2V1000-5	32b data/address optional DES	Internet appliance, industrial control
LavaCORE Configurable Java Processor Core	Derivation	AllianceCORE	V-II	V			38%	20	XC2V1000-5	32b data/address optional DES	Internet appliance, industrial control
Lightfoot 32-bit Java Processor Core	Digital	AllianceCORE	V-II	V	S-II		33%	40	XC2V1000-5	32bit data, 24 bit address, 3 Stage pipeline, Java/C dev. tools	Internet appliance, industrial control, HAVi multimedia, set top boxes
M16450 Universal Asynchronous Receiver Transmitter	Virtual	AllianceCORE				S	29%	60	XC2S50-6	Independently controlled transmit, receive and data interrupts. 16X clock	Serial data applications, modems
M16550A UART with FIFOs	Virtual	AllianceCORE				S	90%	16	XCS20-4	Prog. Data width, parity, stop bits. 16X internal clock, FIFO mode, false start bit detection	Serial data applications, modems
M8254 Programmable Timer	Virtual	AllianceCORE		V	S-II	S					
M8255 Programmable Peripheral Interface	Virtual	AllianceCORE		V	S-II	S	10%	227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable I/O lines, 8-bit bidi data bus	Processor I/O interface
M8259 Programmable Interrupt Controller	Virtual	AllianceCORE				S	28%		XC2S50-6	8 vectored priority interrupts, all 8259/A modes programmable – e.g., special mask, buffer	Real-time interrupt based uP designs
MicroBlaze Soft RISC Processor	Xilinx	LogiCORE	V-II	V	S-II		1%	125	XC2V6000	Soft RISC Processor, small footprint	Networking, communications
OPB Arbiter	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
OPB GPIO	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
OPB Interrupt Controller	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications
OPB Memory Interface (Flash, SRAM)	Xilinx	LogiCORE	V-II	V	S-II			125	Virtex-II	Bundled in the MicroBlaze Development Kit	Processor applications

Xilinx Global Services

Xilinx Global Services

Xilinx Design Services

- Provide extensive FPGA hardware and embedded software design experience backed by industry recognized experts and resources to solve even the most complex design challenge.
- **System Architecture Consulting** — Provide engineering services to define system architecture and partitioning for design specification.
- **Custom Design Solutions** — Project designed, verified, and delivered to mutually agreed upon design specifications.
- **IP Core Development, Optimization, Integration, Modification, and Verification** — Modify, integrate, and optimize customer intellectual property or third party cores to work with Xilinx technology. Develop customer-required special features to Xilinx IP cores or third party cores. Perform integration, optimization, and verification of IP cores in Xilinx technology.
- **Embedded Software** — Develop complex embedded software with real-time constraints, using hardware/software co-design techniques.
- **Conversions** — Convert ASIC designs and other FPGAs to Xilinx technology and devices.

Xilinx Platinum Technical Services

- Senior Applications Engineers
- Dedicated Toll Free Number*
- Priority Case Resolution
- Proactive Status Updates
- Ten Education Credits
- Electronic Newsletter
- Formal Escalation Process
- Service Packs and Software Updates
- Application Engineer to Customer Ratio, 2x Gold Level

*Available in US only

Part Number	Product Description	Duration	Availability
Educational Services			
FPGA13000-4-ILT	Fundamentals of FPGA Design (v4.x)	8	Now
FPGA16000-4-ILT	ISE Design Entry (v4.x) — Instructor-Led course	8	Now
FPGA23000-4-ILT	Designing for Performance (v4.x)	16	Now
LANG1100-3-ILT	VHDL — Introduction to VHDL (v3.x) — Instructor-Led Course	24	Now
LANG11000-4-ILT	VHDL — Introduction to VHDL (v4.x) — Instructor-Led Course	24	November
LANG12000-4-ILT	VERILOG — Introduction to Verilog (v4.x) — Instructor-Led Course	24	Now
PCI1800-3-ILT	PCI CORE Basics (v3.x) — Instructor-Led Course	8	Now
PCI18000-4-ILT	PCI CORE Basics (v4.x) — Instructor-Led Course	8	December
PCI2800-3-ILT	PCI — Designing a PCI System (v3.x) — Instructor-Led Course	16	Now
PC128000-4-ILT	PCI — Designing a PCI System (v4.x) — Instructor-Led Course	16	December
PCI2900-3-ILT	PCI-X — Designing for PCI-X (v3.x)	16	Now
PCI129000-4-ILT	PCI-X — Designing for PCI-X (v4.x)	16	December
ASIC1500-3-ILT	FPGA Design for the ASIC User (v3.x) — Instructor-Led Course	24	Now
ASIC25000-4-ILT	FPGA Design for the ASIC User (v4.x) — Instructor-Led Course	24	November
DSP2000-3-ILT	DSP Implementation Techniques for Xilinx FPGAs	24	November
PROMO-5002-4-LEL	E-Series II (v4)	10	October
TC-1CR	One Training Credit	1	Now
TC-OS-1DY	One Day of On-Site Training	8	Now
TC-OS-2DY	Two days of On-Site Training	16	Now
TC-OS-3DY	Three days of On-Site Training	24	Now
TC-OS-4DY	Four days of On-Site Training	32	Now
TC-OS-5DY	Five days of On-Site Training	40	Now
Platinum Technical Services			
SC-PLAT-SVC-10	1 Seat Platinum Technical Service w/10 education credits		Now
SC-PLAT-SITE-50	Platinum Technical Service site license up to 50 customers		Now
SC-PLAT-SITE-100	Platinum Technical Service site license for 51-100 customers		Now
SC-PLAT-SITE-150	Platinum Technical Service site license for 101-150 customers		Now
Design Services			
DC-DES-SERV	Design Services Contract		

Call Education Services

North America: 877-XLX-CLASS (877-959-2527)
http://support.xilinx.com/support/training/training.htm

Europe: +44-870-7350-548
eurotraining@xilinx.com

Japan: +81-03-5321-7750
http://support.xilinx.co.jp/support/education-home.htm

Asia Pacific: http://support.xilinx.com/support/education-home.htm

Call Design Services

North America & Asia:
Richard Foder: 408-626-4256

Europe:
Alex Hillier: +44-870-7350-516
Martina Finnerty: +353-1-4032469

Call Technical Services

North America: 800-255-7778 or 408-879-5199
http://support.xilinx.com

United Kingdom: +44-870-7350-610
eurosupport@xilinx.com

Japan: +81-03-5321-7750
jhotline@xilinx.com

Other Locations:
http://support.xilinx.com/support/services/contact_info.htm

FREE!
Second Generation XtremeDSP CD Evaluation Kit Available Now!
www.xilinx.com/dsp

The kit includes:

- Simulink® System Modeling Tool from The MathWorks
- System Generator for DSP and ISE 4.1i Development Software from Xilinx®
- FPGA Advantage® from Mentor Graphics®
- SynplifyPro™ from Synplicity®

FORTUNE® 2002
100 BEST COMPANIES TO WORK FOR

XILINX®
The Programmable Logic Company℠

FREE ADMISSION! REGISTER NOW FOR THE CITY NEAREST YOU!

Xilinx, IBM, Wind River Systems, CMP (publisher of EETimes), and other industry leaders invite you to Programmable World 2002. It's an unprecedented event featuring visionary presentations and hands-on technical training sessions from industry leaders, delivering all you need to know about the exciting new era of systems design.

Why You Should Attend

The Forum will highlight new technology poised to revolutionize the way electronic systems are designed.

You will learn how to:

- Implement extreme processing using the industry's fastest FPGA fabric and 300 MHz PowerPC processors
- Design 40+ Gbps applications with next-generation 3.125 Gbps serial I/O technology
- Implement programmable TeraMAC/second DSP systems
- Optimize system performance through dynamic hardware and software (re-)partitioning
- Debug hardware and software in parallel, in your system—at full speed
- Maximize system performance while reducing total system cost

Who Should Attend

This industry event is a must for product development managers, system architects, hardware and software engineers, embedded and DSP engineers, as well as executives, industry analysts, industry press and academia.

Simulcast LIVE Across North America and Europe

Programmable World 2002 will be simulcast live to over 50 cities across North America and Europe including Boston, Dallas, Los Angeles, Ottawa, Raleigh, San Jose, and Paris. Register today to reserve your seat in the city nearest you!

Seating is limited. Register now at www.xilinx.com/pw2002 or visit us in Booth #4644

IBM® XILINX® The Programmable Logic Company℠ WindRiver® CMP United Business Media Electronics Group

PN 0010634